1.0 Introduction

The increasing complexity of satellite systems poses significant challenges to radiation hardness assurance for these systems. Not only are satellites flying commercial parts, fabricated in submicron technologies and featuring novel materials and device types, they are flying these parts in ever greater numbers. At the same time, actual margins on flight applications are being eroded by the use of more powerful and accurate, physics-based radiation transport, circuit simulation and other modeling codes. As complexity increases, space systems may become vulnerable to "rare" failure modes that are difficult to detect with traditional radiation testing methods. An example of such an error was the anomaly seen on the Hubble Space Telescope's solid-state recorder (SSR).[1] In the case of this anomaly, some parts exhibited greater susceptibility to a proton-induced single-event functional interrupt (SEFI) mode than would have been expected from pre-flight testing. Such "surprise" anomalies are disturbing because they may represent failures of traditional RHA methods, or at the very least, failures in their application. Understanding why such surprises occur is potentially crucial—especially given the rapid development of commercial electronics.

To understand why such errors occur—and why to date they have not occurred more often—it is necessary to consider the kinds of errors that affect RHA methodologies. These errors can be viewed as falling into three categories: Poisson errors in SEE, random sampling errors, and systematic errors that result from analysis (or sometimes testing) procedures. Because of the discrete nature of SEE, SEE analysis methods are susceptible to Poisson errors in event counts that go into cross section determinations. The numbers of SEE that occur on orbit are also susceptible to Poisson fluctuations. In contrast, degradation mechanisms, such as TID and displacement damage, are gradual processes, and so not susceptible at present to Poisson errors. (It should be noted that as device sizes shrink they may reach dimensions where a single ion strike may measurably degrade the part. At this point, degradation will have both discrete and gradual components, and Poisson fluctuations will become important.) However, at least in principle, all radiation effects are affected by part-to-part variability. The process of bounding this variability is subject to random sampling errors and potentially to systematic errors.

In conventional RHA methodologies, parts in the flight lot are viewed as being drawn from some larger "parent" population of parts—perhaps the wafer diffusion lot or an even larger ensemble. Parts drawn from the parent population for a flight lot have a probability of failure, $p(S)$ when subjected to a stress, S, such as a TID, proton flux, etc. RHA generally relies on inference of $p(S)$ from test results for a (usually small) representative Radiation Lot Assurance Test (RLAT) or Radiation Characterization (RC) sample drawn from the same parent population. As such, RHA methods are susceptible to random sampling errors, which follow binomial statistics. While such random sampling errors decrease with sample size (see figure 1), economic constraints usually restrict samples to sizes that, according to binomial statistics, still leave considerable uncertainty about flight-lot hardness. One of the reasons why confidence increases only slowly with sample size is that binomial statistics make no assumptions about details of the failure distribution, instead using pass/fail test results to infer the relative proportions of good and bad parts for a given applied stress. In order to draw stronger conclusions, most methods make assumptions—perhaps motivated by experience, empirical evidence, device technology or failure mechanism models—about the characteristics of the failure distribution.

In general, the assumptions that must be made are not particularly stringent. For example, if the distribution is not bimodal, then increasing the radiation design margin (RDM—the ratio of the failure dose to the application dose) can significantly increase confidence in the proportion of good parts. Unfortunately, to the degree that the actual failure distribution differs from that assumed, such procedures risk introducing systematic error. RHA methods may draw erroneous conclusions if they fail to account accurately for such errors. (See figure 2.)

At present, such sampling and systematic errors are more of a concern for degradation mechanisms than for SEE. However, in the future it is likely that inaccuracies in current SEE calculation methods will motivate development of more accurate, physics-based SEE rate calculation methods. For these more accurate methods, part-to-part variability will likely be of greater concern.

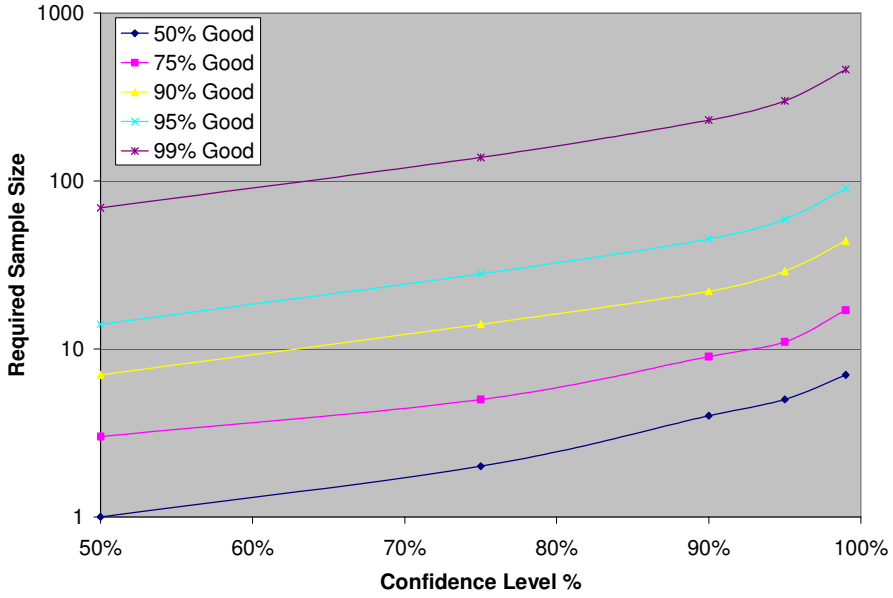

Figure 1 Sample size required to achieve a given confidence that the given percentage of parts is good.

In this note we consider the three error types mentioned above. Because Poisson errors pertain mainly to SEE and because they often dominate contributions due to part-to-part variability, we begin by considering the possible effects of Poisson fluctuations on SEE rate predictions, the robustness of SEE rate calculation methods to such errors and the systematic errors characteristic of SEE rate calculation algorithms. We then consider the further complication of SEE induced by the spallation products of proton-nuclear interactions.

Next, we examine the implications of sampling errors and of the tradeoff between these random errors and the potential for systematic errors implicit in making assumptions about the parent population failure distribution. Previous treatments of statistical inference in RHA have typically assumed that parent population failure distributions are normal or lognormal.[2],[3],[4],[5] However, these treatments have focused primarily on small flight lots (of order 10 parts), where consequences of departures from the assumed distribution are usually negligible. In addition, at the time these studies were published, traceability of parts to a particular wafer—or even a particular region within a wafer—was much more common that it is today. At present, we are confronted with architectures that are using hundreds or even thousands of "identical" parts—and some of these parts are purely commercial, fabricated with new techniques and materials and with limited traceability. For such complicated systems, the effects of any random or systematic errors may be greatly magnified. As such, we seek to develop a self-consistent treatment not just of sampling errors, but also of any systematic errors introduced by assuming a given form for the failure distribution. To do this, we must first develop a method for constraining the types of distributions that must be considered.

In general, the constraints that must be introduced are not stringent—mainly the elimination of distributions with multiple modes and/or long, thick tails (e.g. the Cauchy distribution in figure 2). Where possible, we do this using archival qualification and RLAT data. However, if such data do not exist (the usual case), we examine whether some aspects of device technology and of radiation effects basic mechanisms may be used to at least increase confidence that the failure distribution is well behaved.

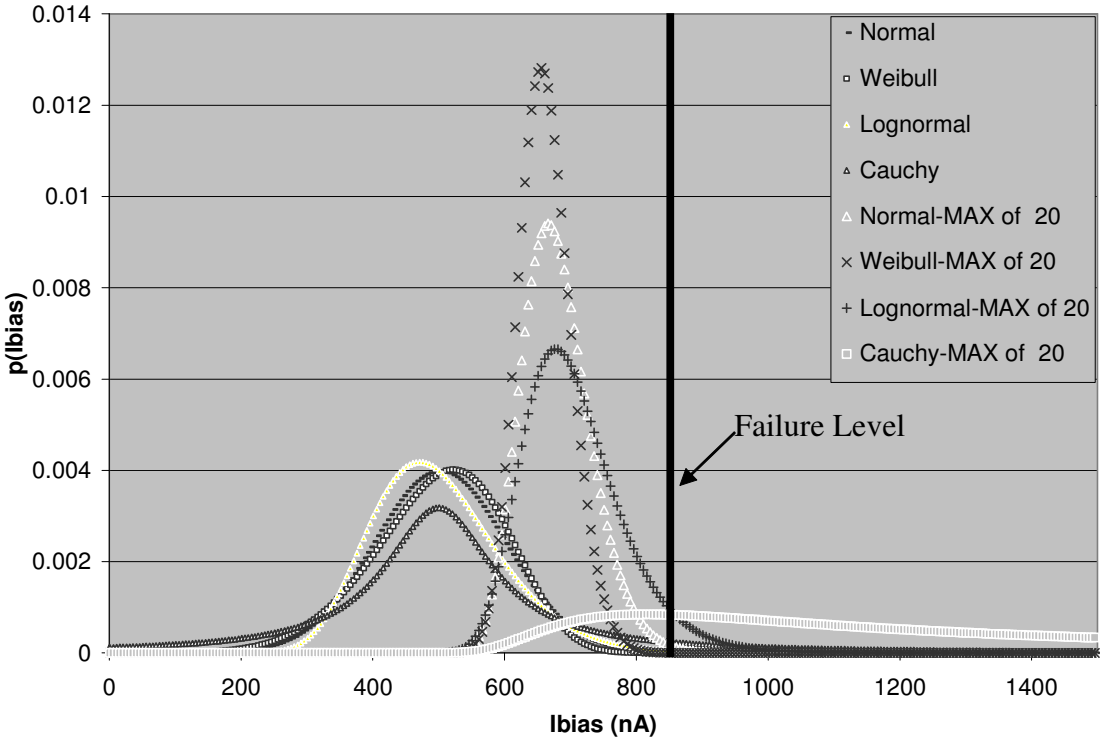

Figure 2 Radiation data samples are typically too small to distinguish between distribution types. The data shown are consistent with a Normal distribution (μ=500, σ=100), a Weibull (s=5.8, w=540) a lognormal (μln=6.19, σ=0.198) or even a Cauchy (mode=500, width=100). Yet, the distributions for the worst-case leakage current can vary dramatically for flight lots as small as 20 parts. The distribution of the maximum assuming a Cauchy distribution is nearly flat with I_{bias}.

Having eliminated from consideration the most problematic distribution forms, we then consider the implications of the failure distribution having each of three different forms that can be motivated by physical or mathematical arguments: Normal, Lognormal and Weibull. In addition to being physically or otherwise reasonable, these three distributions also possess some other interesting traits. The Normal distribution is symmetric about the mean, and possesses zero excess kurtosis. The Lognormal distribution has a slightly thicker right tail (positive skewness) than the Normal distribution, while the Weibull (except when very broad) has a slightly thicker left tail (negative skewness). As a result, simultaneously fitting data to all three distributions provides a measure of the systematic errors associated with assumption of a particular distribution form—and particularly of the sensitivity of the conclusions of an analysis to higher moments of the distribution that a small sample is unlikely to measure. In what follows, we treat part-to-part variation as it applies to degradation mechanisms—particularly TID degradation. We also present the probability distribution of interest as a probability of observing a particular amount of degradation for a given TID (as in figure 2 above), rather than as a probability of failure at a given dose. However, the conclusions we arrive at are easily adaptable to the latter presentation. If part-to-part variability is important for a particular SEE analysis, it can also be readily adapted and combined with the conclusions for the effects of Poisson errors—which we now consider.

2.0 SEE as Poisson Processes

Most single-event analyses assume that SEE occur at a constant mean rate, neglecting part-to-part variability, modulations in the galactic cosmic ray environment and short-term fluctuations in the solar heavy-ion environment.[7] At most, different mean rates will be estimated for conditions such as

Solar Maximum, Solar Minimum and Solar Events.[6] Under these circumstances, SEE occurrence varies with Poisson statistics with mean rate $\mu=\int F(LET)\bullet\sigma(LET)dLET$, where *F(LET)* is the fluence of ions of a given LET capable of causing the given effect and $\sigma(LET)$ is the effective vulnerable area for the effect. The probability of the same fluence causing *n* rather than μ SEE is

$$P(n)=\frac{\mu^n e^{-\mu}}{n!} \qquad (1)$$

2.1 Confidence Intervals for cross section vs. LET

Equation (1) can be used to bound the cross section for a particular SEE at a given confidence level. For example, a null result for an SEE at a particular LET could represent a downward fluctuation from a mean susceptibility, rather than an indication that the device is not susceptible. Assuming Poisson statistics, there is a 5% chance of seeing 0 events for distributions with mean ~2.997, implying that if 0 events are seen for a fluence *F*, then at the 95% CL the mean corresponding to the observation must be less than 2.997, and that the limiting cross section $\sigma<2.997/F$. The process of deriving a 95% CL from an observed cross section vs. LET curve is a bit more involved. However, by using likelihood ratios based on the hypotheses that 1) the observed count is the mean, or 2)the observed count represents a downward fluctuation from the mean. The logarithm of the likelihood ratio scales as the χ^2 distribution for the numbers of points allowed to vary

$$\ln\left(L_2/L_1\right)=-0.5\chi^2(1-CL, DOF=\#\ pts\ ) \qquad (2)$$

Although there are many different curves that can be constructed at the same confidence level, we are interested in the curve that generates the worst case rate for that CL. Typically, the highest rate will be obtained by allowing the lowest-LET points to fluctuate, since these points are usually based on the lowest counts (and so have the largest Poisson errors) and coincide with the higher fluxes of galactic cosmic rays. In this case, one is on the rising edge (i.e. below the knee) of the Weibull, so an additional constraint can be imposed—that cross sections for low LET should be lower than those for high LET. Thus if we had the cross section vs. LET curve shown in the first two columns of Table I, the event counts in the third column have a combined probability of ~0.5. Note that the Poisson probability for observing one count less than the mean is the same as observing the mean, so we increase the fourth, fifth and higher bins by one.

Table I: Sample Event # vs LET and Fluctuation

LET	Observed Events	Inferred Mean	Probability
2.7	0	2	
3.1	1	3.5	
3.8	2	4.5	
5.4	10	11	~0.05
7.8	15	16	
• • •	• • •	• • •	

.

Provided the σ vs. LET curve is of sufficient quality (i.e. small statistical errors and a reasonable number of data points), a very similar result can be obtained by using the worst-

case shape and width parameters consistent with the confidence contour CL of the Maximum Likelihood Estimators. (See below in section 2.3.)

2.2 Defining a Test Fluence

One can also use Poisson statistics to develop guidelines for testing parts for "rare" effects. In the most simplistic view of SEE, an error occurs when an ion of sufficient LET strikes a vulnerable feature in the chip (e. g., the drains of vulnerable transistors in a CMOS register). The ion strikes can be viewed as cylinders of charge, with the charge density having a complicated radial and temporal dependence. If the vulnerable areas in the cell are larger than the track radius, then the radial dependence of the charge track is irrelevant to the SEE mechanism. Then the question is whether each vulnerable area has been probed with the desired level of confidence, *CL*, by a flux, *F*, of sufficiently high LET ions. This means that the probability that a given minimum vulnerable area, *A*, has not been probed is 1-*CL*, so the required ion flux is

$$F = \frac{-\ln(1-CL)}{A} \qquad (3)$$

Thus, a fluence of 10^7 ions per cm^2 would be sufficient to probe vulnerable feature areas down to 36 square microns with 95% confidence.

On the other hand, if the minimum vulnerable feature sizes on the device are small compared to the track size, one may choose to consider the area in the denominator of (2) to be the cross sectional area of the charge track on the chip. (See Figure 3.) At this scale, one must consider track structure effects—particularly the radial charge distribution and its time evolution. Device structures (e.g. isolation), characteristics (e.g. doping) and operating conditions (e.g. speed, biases, etc.) that influence charge collection may also be important. Fortunately, computer simulations and experimental comparisons of ion-induced and laser-induced charge distributions have indicated that as long as the evolution of the charge distribution (typically picoseconds) is rapid compared to the timescales of charge collection and device operation (typically nanoseconds in many current technologies), the initial charge distribution is of limited importance.[7],[8],[Kobayashi] In this case, the relevant value for *A* is on the order of 10 square microns—implying that a fluence on the order of 1E7-1E8 ions per cm^2 is adequate to illuminate rare effects at the 95% CL.

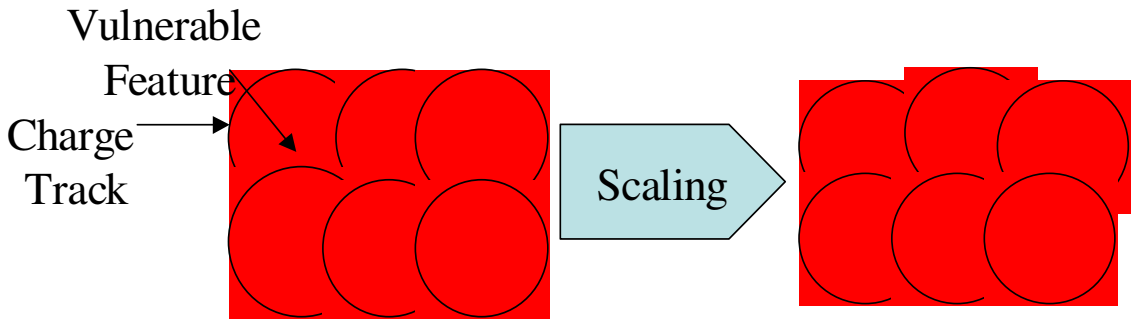

Figure 3 Although scaling increases the density of cells on a microcircuit, the finite size of charge columns due to ion tracks means that the ion fluence needed to probe all features on an IC does not increase commensurately. However, for deep submicron technologies, track structure or circuit effects may become important.

As device sizes decrease below 100 nm, track structure effects will become increasingly important, with the effects on transistors depending on where they are within the ion-strike charge distribution. Indeed, since one ion strike may affect hundreds of transistors at this scale, SEE analysis may require consideration of circuit-level as well as the device-level effects. Moreover, for such device dimensions the concept of particle LET may break down, with energy deposition itself being a Poisson process. Even TID and displacement

damage may become stochastic processes at these scales, and a treatment that better reflects the physics of these processes will be necessary.[10]

At present, the fluence needed to probe parts to the limiting feature size is in the 1E7-1E8 cm^{-2} range. This range will reveal vulnerabilities with cross sections as low 3E-7 – 3E-8 cm2 at the 95% CL. For the upper end of this fluence range, it may be necessary to irradiate more than one part per ion to avoid TID degradation during the run.

The above considerations pertain to determination of device susceptibility to single-event effects. In addition, SEE hardness assurance also seeks to predict—or at least bound—the on-orbit SEE rate. SEE rate calculation typically involves fitting the SEE σ vs. LET curve to a cumulative Weibull distribution function. One then convolutes this device response with the galactic cosmic ray spectrum over the distribution of chords in a rectangular-parallelepiped charge collection volume to yield the SEE rate. Often, because of the uncertainties in SEE rate calculation, the fit is done conservatively by eye, rather than by more sophisticated methods. While this method is effective for producing bounding rates, the lack of confidence intervals makes it difficult to compare different estimates for the same device and to incorporate SEE rate estimates in reliability calculations.

2.3 Analysis of SEE data

Although most SEE calculation algorithms assume the form of the cross section vs. LET curve is Weibull, this decision is motivated more by the versatility of the Weibull than by physics. Indeed, Edmonds[??] has suggested that a Lognormal distribution does a better job accounting for the effects of diffusion due to its somewhat thicker tail. Because most available rate calculation methods are tailored to the Weibull fit, we will confine our discussion to it, with the implicit assumption that almost all of our considerations apply to lognormal fits as well.

2.3.1 Fitting SEE Data and Confidence Intervals

Although SEE data can be fit to the desired function by a variety of means, Maximum Likelihood Estimation (MLE) has the advantage that it provides not only the fit parameters but also confidence intervals for them. However, the maximum likelihood method usually fits data to a probability density function, while SEE data are usually fit to a cumulative Weibull. The way around this is to express the data in terms of the differential increase in cross section from one LET step to the next. (This is most easily done by normalizing all event counts to a single fluence, say 10^7 ions/cm^2.) This interval data can then be fit using standard MLE procedures for the Weibull. One can then explore the parameter space consistent with the desired CL to determine which parameters yield the worst-case rate for that CL. An example of such a contour is shown in figure 4, which shows the cross section vs. LET curve for transients in the Analog Devices AD623 op amp. Worst-case rates will generally be associated with small width and shape parameters. One feature of this method is that downward fluctuations result in negative counts, so that interval cannot be used. However, such downward fluctuations will usually occur near or above saturation, and will not strongly influence the fit in any case. Because the data giving rise to figure 4 had excellent statistics (>1000 transients for each LET value), the calculated rates differ by only 20% due to the fit even at the 99% CL. Had the data been based on counts 10x smaller (still good statistics), the uncertainty becomes ±~50% at the 99% CL (or about 35% at the 95% CL). A further factor of 10 in the numbers of counts yields a rate that could be off by a factor of 3 at the 99% CL (>2x at the 95% CL).

S \ W	28.3	28.38	28.46	28.54	28.62	28.7	28.78	28.86	28.94	29.02	29.1	29.18	29.26	29.34	29.42	29.5	29.58	29.66	29.74	29.82	29.9	30	30.1	30.1
1.48	11.95	11.04	10.25	9.58	9.03	8.60	8.29	8.10	8.02	8.06	8.21	8.47	8.85	9.33	9.91	10.61	11.41	12.31	13.31	14.42	####	####	####	####
1.485	11.07	10.11	9.29	8.58	8.00	7.54	7.19	6.97	6.86	6.87	6.99	7.22	7.56	8.02	8.58	9.25	10.02	10.90	11.88	12.97	####	####	####	####
1.49	10.29	9.30	8.43	7.69	7.07	6.57	6.20	5.94	5.80	5.78	5.87	6.07	6.39	6.81	7.35	7.99	8.74	9.59	10.55	11.62	####	####	####	####
1.495	9.62	8.58	7.68	6.90	6.25	5.72	5.30	5.01	4.84	4.79	4.85	5.02	5.31	5.71	6.21	6.83	7.56	8.39	9.32	10.36	####	####	####	####
1.5	9.05	7.98	7.03	6.22	5.53	4.96	4.51	4.19	3.99	3.90	3.93	4.07	4.33	4.70	5.18	5.77	6.47	7.28	8.19	9.21	####	####	####	####
1.505	8.58	7.47	6.49	5.64	4.91	4.31	3.83	3.47	3.23	3.11	3.11	3.23	3.45	3.80	4.25	4.82	5.49	6.27	7.16	8.15	9.25	####	####	####
1.51	8.22	7.07	6.05	5.16	4.39	3.76	3.24	2.85	2.58	2.43	2.39	2.48	2.68	2.99	3.42	3.96	4.61	5.36	6.23	7.20	8.27	9.45	####	####
1.515	7.97	6.77	5.71	4.78	3.98	3.31	2.76	2.33	2.02	1.84	1.78	1.83	2.00	2.29	2.69	3.20	3.82	4.55	5.39	6.34	7.39	8.55	9.81	####
1.52	7.81	6.58	5.48	4.51	3.67	2.96	2.37	1.91	1.57	1.36	1.26	1.28	1.42	1.68	2.05	2.53	3.13	3.84	4.65	5.58	6.61	7.74	8.98	####
1.525	7.76	6.49	5.35	4.34	3.46	2.71	2.09	1.59	1.22	0.97	0.84	0.83	0.94	1.17	1.51	1.97	2.54	3.22	4.01	4.91	5.92	7.03	8.25	9.57
1.53	7.81	6.49	5.31	4.26	3.35	2.56	1.90	1.37	0.97	0.68	0.52	0.48	0.56	0.76	1.07	1.50	2.04	2.70	3.47	4.34	5.32	6.41	7.61	8.91
1.535	7.96	6.60	5.38	4.29	3.34	2.51	1.82	1.25	0.81	0.49	0.30	0.23	0.28	0.45	0.73	1.13	1.65	2.28	3.02	3.87	4.83	5.89	7.07	8.35
1.54	8.21	6.81	5.55	4.42	3.43	2.56	1.83	1.23	0.75	0.40	0.18	0.07	0.09	0.23	0.48	0.86	1.34	1.95	2.66	3.49	4.42	5.47	6.62	7.88
1.545	8.56	7.12	5.82	4.65	3.61	2.71	1.94	1.30	0.79	0.41	0.15	0.01	0.00	0.11	0.33	0.68	1.14	1.71	2.40	3.20	4.11	5.14	6.27	7.50
1.55	9.02	7.53	6.18	4.97	3.90	2.96	2.15	1.48	0.93	0.51	0.22	0.05	0.01	0.08	0.28	0.59	1.03	1.57	2.24	3.01	3.90	4.90	6.01	7.22
1.555	9.57	8.04	6.65	5.40	4.29	3.31	2.46	1.75	1.17	0.71	0.39	0.18	0.11	0.15	0.32	0.60	1.01	1.53	2.17	2.92	3.78	4.75	5.84	7.03
1.56	10.22	8.65	7.21	5.92	4.77	3.75	2.87	2.12	1.50	1.01	0.65	0.41	0.30	0.32	0.45	0.71	1.08	1.58	2.19	2.91	3.75	4.70	5.76	6.93
1.565	10.97	9.35	7.88	6.54	5.35	4.29	3.37	2.58	1.92	1.40	1.00	0.74	0.59	0.58	0.68	0.91	1.26	1.72	2.30	3.00	3.81	4.74	5.78	6.93
1.57	11.81	10.15	8.64	7.26	6.02	4.93	3.97	3.14	2.45	1.89	1.46	1.15	0.98	0.93	1.00	1.20	1.52	1.96	2.51	3.18	3.97	4.87	5.89	7.01
1.575	12.76	11.05	9.49	8.07	6.80	5.66	4.66	3.79	3.06	2.47	2.00	1.67	1.46	1.38	1.42	1.59	1.87	2.28	2.81	3.46	4.22	5.10	6.09	7.19
1.58	13.80	12.05	10.44	8.98	7.66	6.49	5.45	4.54	3.78	3.14	2.64	2.27	2.03	1.92	1.93	2.06	2.32	2.70	3.20	3.82	4.56	5.41	6.37	7.45
1.585	14.94	13.14	11.49	9.99	8.63	7.41	6.33	5.39	4.58	3.91	3.38	2.97	2.70	2.55	2.53	2.63	2.86	3.21	3.69	4.28	4.99	5.81	6.75	7.81
1.59	16.18	14.33	12.64	11.09	9.69	8.43	7.31	6.33	5.48	4.78	4.20	3.76	3.45	3.27	3.22	3.30	3.49	3.82	4.26	4.82	5.51	6.31	7.22	8.26
1.595	17.51	15.62	13.88	12.29	10.84	9.54	8.38	7.36	6.48	5.73	5.12	4.65	4.30	4.09	4.00	4.05	4.22	4.51	4.92	5.46	6.12	6.89	7.78	8.79
1.6	18.93	17.00	15.21	13.58	12.09	10.74	9.54	8.48	7.56	6.78	6.13	5.62	5.24	5.00	4.88	4.89	5.03	5.29	5.68	6.19	6.82	7.56	8.43	9.41
1.605	20.46	18.47	16.64	14.96	13.43	12.04	10.80	9.70	8.74	7.92	7.24	6.69	6.28	6.00	5.85	5.82	5.93	6.16	6.52	7.00	7.60	8.33	9.17	####
1.61	22.07	20.04	18.17	16.44	14.86	13.43	12.15	11.01	10.01	9.15	8.43	7.85	7.40	7.09	6.90	6.85	6.92	7.13	7.45	7.91	8.48	9.18	9.99	####

Figure 4 Maximum likelihood estimation yield confidence ellipses as well as estimators for the distribution. In this figure, the values correspond to the negative change in the log likelihood. The magenta cell indicates the best-fit parameters (S=1.545, W=29.26). The inner boundary of the yellow ring corresponds to the 90% confidence intervals, the outer boundary to the 95% CL. Similarly, the inner boundary of the turquoise ring corresponds to 99% CL and the outer boundary to the 99.9% CL. The SEE rate for this confidence level corresponds to the parameters on the ellipse that yield the highest rate.

2.3.2 Robustness of SEE Rate Calculations

As mentioned briefly, SEE rate calculations are relatively robust against minor part-to-part variations. In part, this is because random and systematic errors inherent in the methods render SEE rates uncertain by a factor of two at best. In addition, the methods are relatively forgiving of small errors in the parameters that determine the rate. This is most easily seen for the somewhat less precise Figure of Merit (FOM) rate,[11],[12] which estimates the rate in terms of the limiting cross section, σ_{lim}, and the threshold LET, LET_{th}, (usually the LET where the cross section reaches a quarter of its saturated value)

$$R \approx C \times \sigma_{\text{lim}} \Big/ LET_{th}^2 \quad , \tag{4}$$

where C is a constant. Ignoring any dependence of σ_{lim} on LET, we see that the error on the rate can be expressed as

$$dR \approx C / LET_{th}^2 \times \left[\left(d\sigma_{\text{lim}} \right) - \left(\frac{2\sigma_{\text{lim}} \times dLET_{th}}{LET_{th}} \right) \right] = R \times \left[\left(\frac{d\sigma_{\text{lim}}}{\sigma_{\text{lim}}} \right) - 2 \times \frac{dLET_{th}}{LET_{th}} \right] \tag{5}$$

Thus, the rate changes only logarithmically in errors on σ_{lim} and LET_{th}. In general, the bracketed quantities in the rightmost version of equation 5 will be small, suggesting limited variability in the rate compared to the inherent 2x uncertainty in rate calculation methods. While current SEE rate calculation methods are considerably more complicated, they are similarly forgiving. Indeed, as seen above, provided the cross section vs. LET curve has enough points and that the statistical uncertainty on the points is sufficiently small, the WC 95% CL rate will differ little from the best-fit rate.

In addition, on-orbit rates are also susceptible to Poisson fluctuations, so until sufficient statistics (i.e. SEUs, etc.) have been accumulated, it can be difficult to estimate the on-orbit rate with any confidence. Since most applications demand

relatively low susceptibility to SEE, it is rare to be able to state with confidence that an on-orbit rate is different from that predicted.

2.4 Proton-Induced SEE

If a device exhibits sensitivity to an SEE mode for LET≤15 MeVcm^2/mg, proton-induced SEE mechanisms may introduce yet another layer of complexity—the interaction of the proton with the Si nucleus and the subsequent charge generation by the products of that interaction. In general, one in 10^4-10^6 protons will produce recoil products with sufficient LET and range to cause an SEE, so proton-induced SEE cross sections will be lower by a corresponding factor relative to the heavy-ion induced mode. Thus, probing the microcircuit to the same level as for heavy ions could require proton fluences of order 10^{13} protons/cm^2. While one can argue that only a small minority of missions would be likely to encounter proton fluences this large, flying hundreds or thousands of a given part significantly increases the vulnerable area. Since the proton LET is about 1000 times less than that of a heavy ion, TID degradation of parts in a proton test may require multiple parts to even observe the effect. Proton energy may also be important, since the ranges of recoil products increase with proton energy. Moreover, Reed et al. recently emphasized the importance of recoil directionality and proton angle of incidence for some thin device geometries (e.g. SOI).[13] And Buchner et al.[14] recently showed that similar angular effects can produce enhancements in multibit upsets in DRAMs. Given all of these variables, it is not surprising that some technologies exhibit sensitivity to an SEE at low LET and fail to exhibit sensitivity in subsequent proton testing. In general, such failures to observe proton-induced SEE are an indication that these effects would be rare on orbit. However, if the flight application involves large numbers of parts, the probability of observing these effects rises—and if susceptibility to the effect varies from part to part, it may rise even faster than linearly with the number of parts being flown.

Moreover, thorough characterization of proton-induced SEE involves irradiation of parts at several proton energies and in some cases (because of the angular dependence of recoil products), over angles. As such, for most technologies, the device cross section vs. proton energy curve will require irradiation of multiple parts per proton energy. In such cases, one cannot avoid the question of part-to-part variability. This introduces additional complications that rigorous statistical analysis alone cannot resolve economically.

3.0 Treating Part-to-Part Variability

As mentioned above, binomial statistics requires very large sample sizes to infer details of the parent failure distribution with any degree of confidence. This forces RHA methods to make assumptions about the form of the failure distribution. On the one hand, it is quite reasonable to assume that the failure distribution is not completely arbitrary for semiconductor devices—the fabrication of which ranks among the most carefully controlled processes ever undertaken by humankind. On the other hand, erroneous assumptions introduce systematic errors into the process, potentially leading to surprise failures. Some parts are known to exhibit bimodality in TID response, both lot-to-lot and in some cases, intra-lot. If a part's radiation response distribution were to exhibit thick tails, this would also be problematic. Very large flight lots are much more likely to sample from multiple modes or thick tails of a distribution than are small RLAT or other test lots. As such, it can be risky to base results on an assumed distribution for small test lots. The question is what distribution forms are reasonable to assume, and can the types of distributions be motivated or restricted. Here we propose three distribution types that are physically reasonable under various scenarios and which allow exploration of the sensitivity of analysis results to factors such as distribution symmetry (or skewness) and the characteristics of distribution tails. The distributions we propose are the normal, lognormal and Weibull, and the rationales, advantages and disadvantages of each are as follows.[15],

Normal Distribution: There are many reasons for making an initial assumption of normality for the failure distribution—even over and above its convenience and familiarity. The sample mean and standard deviations have good convergence properties as estimators of their population counterparts. The Central Limit Theorem guarantees that at least near the mode most distributions will look more or less Normal. The small sample size behavior is well known, being described by the Student's t distribution. Also, the Normal distribution is its own

conjugate prior (under some circumstances) in Bayesian statistics, so a Bayesian (or quasi-Bayesian) analysis may be performed. A Bayesian approach has much to recommend it for RHA, since it is tailored for situations in which decisions must be made on the basis of sparse information. Moreover, regardless of the values of distribution parameters (μ,σ), the distribution remains symmetric (skewness=0) and the proportions of the distribution in the peak and the tails remains constant (excess kurtosis=0).

Unfortunately, the normal distribution is not physical for many situations, since it is defined for both positive and negative stress. The fact that most distributions appear Normal in their central behavior may obscure deviations from normality in the tails. However, as long as the ratio σ/μ remains small, and as long as the flight lot is not very large (~10 or so), the normal distribution is a convenient and not an unreasonable approximation. One significant advantage is that one can use the student's t distribution to estimate the sample size, N, necessary to achieve a given confidence level (*CL*) for a given allowable error on the mean, δQ and population standard deviation, σ (or, if σ is not known, one may substitute the sample standard deviation, *s)*:

$$N = (T^{-1}(CL, DOF)(\frac{\sigma}{\delta Q}))^2 \qquad (6)$$

There are several characteristics of this equation that are noteworthy. One can estimate σ based on archival data if it is available. Alternatively, if one substitutes the sample standard deviation for that of the population, the procedure is more complicated. Since the sample standard deviation is not known *a priori*, one must express δQ in terms of *s*—e.g. by taking δQ=*s (*or *0.5 s…)*, which places the mean within *s* (or 0.5 s) of the sample mean for the desired confidence level. Substitution of *s* for σ also introduces statistical error into the process. Since *s* for a single sample may vary considerably from the population standard deviation σ, it may be desirable to use a more conservative estimate of σ determined from *s* and its confidence intervals. If M_4 is the fourth central moment of the parent distribution, the variance of the variance for a sample of size n can be expressed as:

$$\mathrm{var}(s^2) = (1/n)[M_4 - \frac{n-3}{n-1}\sigma^4] \qquad (7).$$

By substituting the sample moments for the population moments, one can estimate the error on the variance. Alternatively, one can use the value of s for the appropriate confidence interval of the maximum likelihood estimator for σ.

Finally, T^{-1} depends on *N* through the number of degrees of freedom, so this equation is most easily solved by using a table of T^{-1}, increasing N step by step, and looking in the column for (*N-1*) degrees of freedom until one finds a value less than 1-CL. For example, in Figure 5, a sample size of 4 is needed to achieve 90% confidence that the sample mean is within a standard deviation of the real mean. Note that requiring knowledge of the mean to within half the standard deviation to the same confidence level quadruples the required sample size. Although, the assumption of normality is arbitrary and unphysical, it represents a good approximation to the central behavior of many distributions—and for small flight-lot or sample sizes, one is unlikely to be sampling from the extremes in any case. Moreover, equation (6) is extremely valuable for estimating required sample sizes. Even if we know that our failure distribution cannot be Normal, convergence of parametric estimators for many other distributions often behaves in a manner similar to that for the Normal distribution.

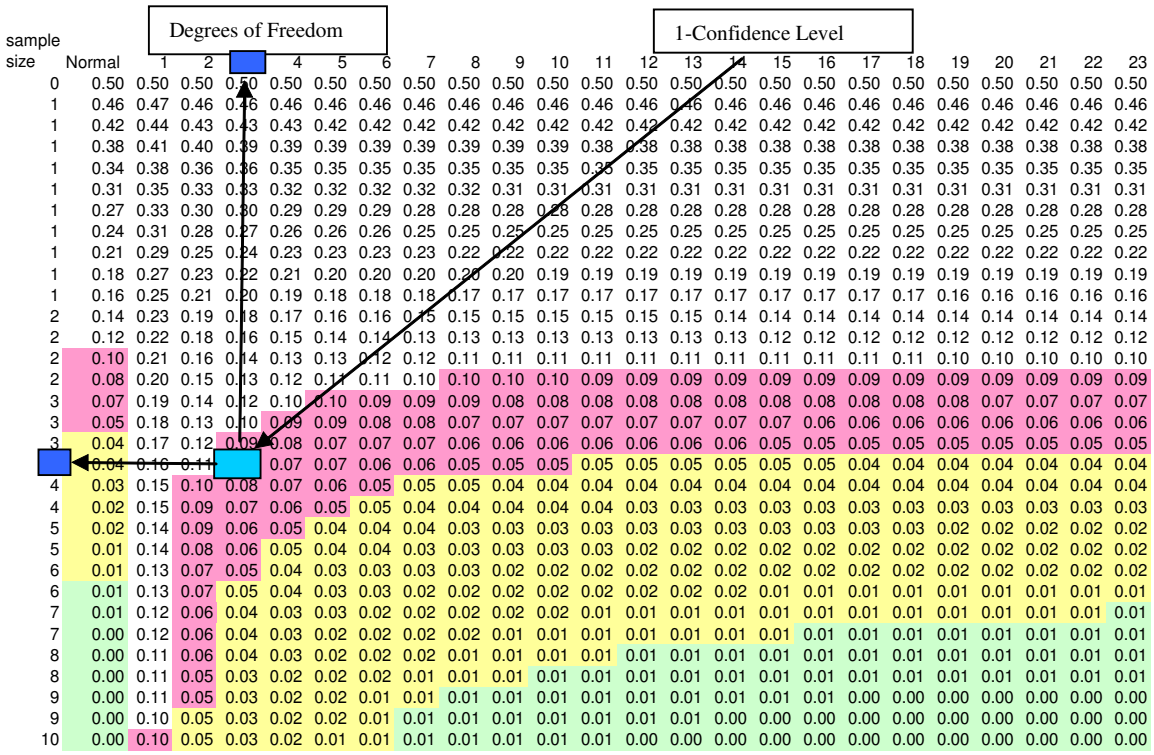

sample size	Normal	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
0	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
1	0.46	0.47	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46	0.46
1	0.42	0.44	0.43	0.43	0.43	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42	0.42
1	0.38	0.41	0.40	0.39	0.39	0.39	0.39	0.39	0.39	0.39	0.39	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38	0.38
1	0.34	0.38	0.36	0.36	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35
1	0.31	0.35	0.33	0.33	0.32	0.32	0.32	0.32	0.32	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31
1	0.27	0.33	0.30	0.30	0.29	0.29	0.29	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28	0.28
1	0.24	0.31	0.28	0.27	0.26	0.26	0.26	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
1	0.21	0.29	0.25	0.24	0.23	0.23	0.23	0.23	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22	0.22
1	0.18	0.27	0.23	0.22	0.21	0.20	0.20	0.20	0.20	0.20	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19	0.19
1	0.16	0.25	0.21	0.20	0.19	0.18	0.18	0.18	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.17	0.16	0.16	0.16	0.16	0.16
2	0.14	0.23	0.19	0.18	0.17	0.16	0.16	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.14
2	0.12	0.22	0.18	0.16	0.15	0.14	0.14	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.12
2	0.10	0.21	0.16	0.14	0.13	0.13	0.12	0.12	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.11	0.10	0.10	0.10	0.10	0.10
2	0.08	0.20	0.15	0.13	0.12	0.11	0.11	0.10	0.10	0.10	0.10	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09	0.09
3	0.07	0.19	0.14	0.12	0.10	0.10	0.09	0.09	0.09	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.07	0.07	0.07	0.07
3	0.05	0.18	0.13	0.10	0.09	0.09	0.08	0.08	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.06
3	0.04	0.17	0.12	0.09	0.08	0.07	0.07	0.07	0.06	0.06	0.06	0.06	0.06	0.06	0.06	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
4	0.04	0.16	0.11	0.08	0.07	0.07	0.06	0.06	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
4	0.03	0.15	0.10	0.08	0.07	0.06	0.05	0.05	0.05	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
4	0.02	0.15	0.09	0.07	0.06	0.05	0.05	0.04	0.04	0.04	0.04	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
5	0.02	0.14	0.09	0.06	0.05	0.04	0.04	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.02	0.02	0.02	0.02	0.02
5	0.01	0.14	0.08	0.06	0.05	0.04	0.04	0.03	0.03	0.03	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
6	0.01	0.13	0.07	0.05	0.04	0.03	0.03	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
6	0.01	0.13	0.07	0.05	0.04	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
7	0.01	0.12	0.06	0.04	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
7	0.00	0.12	0.06	0.04	0.03	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
8	0.00	0.11	0.06	0.04	0.03	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
8	0.00	0.11	0.05	0.03	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
9	0.00	0.11	0.05	0.03	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
9	0.00	0.10	0.05	0.03	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
10	0.00	0.10	0.05	0.03	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Figure 5 Determining the sample size using equation (5) involves looking at the values of the t distribution as a function of confidence level and the degrees of freedom. In this case, the confidence level range of interest –90-95%—is shown in pink. The smallest sample size, n, that yields a self-consistent number of degrees of freedom (n-1) in this range is n=4 (see shaded values).

Lognormal Distribution: While the normal distribution is often not physical, the lognormal distribution is not only physically possible (i.e. defined only for positive stresses), it is also physically reasonable for many systems. In particular, if the incremental damage to a system exposed to a stress is proportional to the damage previously sustained, the failure will follow a lognormal distribution. As might be expected, the fact that the logarithms of the stress are normally distributed means that the convergence properties of the lognormal distribution are often similar to those of the normal distribution. The disadvantages to the lognormal distribution are that it is somewhat more difficult to work with, and there is no analog to the Student's t distribution. As such, the small sample behavior is not as well defined as for the Normal distribution. The equivalent of the student's t distribution can be estimated numerically. However, there is no "standard" lognormal t distribution, since the distribution will vary depending on the lognormal distribution parameters (mean and standard deviation). Unlike the normal, the lognormal distribution is skewed in the positive direction. For values of the lognormal σ likely to occur in RHA (i.e. $\sigma<0.2$) the skewness will be relatively slight. As mentioned above, comparing the results of assuming Normal and Lognormal distributions can test the sensitivity of the results to slight positive skewness.

Weibull Distribution: Like the lognormal distribution, the Weibull distribution has the advantage that it is physically justifiable. Systems in which a variety of failure mechanisms compete (weakest-link failures) tend to fail according to a Weibull distribution. In addition, the Weibull assumes a variety of morphologies—from exponential when the shape parameter, $\beta=1$ to a delta function as $\beta\rightarrow\infty$. As with the lognormal distribution, there is no analog to the Student's t distribution defined for the Weibull, but the small sample behavior can be estimated numerically as a function of the shape parameter. Moreover, for shape parameters

greater than 3.6 (usually, s>5.8 except for very broad distributions), the Weibull is skewed toward the left. Thus, along with the lognormal, the Weibull allows investigation of how sensitive conclusions of an analysis may be to nonzero skewness (left or right).

Although the three distributions considered here may not be exhaustive of all possible failure distributions, they do represent are the main distributions that can be justified on the basis of physics. They also represent a fairly broad range of distribution morphologies that allow investigation of the effects of skewness and slight thickening in the left (for the Weibull) and right (for lognormal) tails relative to the Normal distribution. The choice of one distribution may be based on a variety of criteria. In some cases, there may be a compelling physical reason one expects one distribution form over others. There may be empirical arguments. For instance, if significant amounts of archival data are available, and if the inter-lot and intra-lot spreads are comparable (see figure 6), one may be able to combine multiple lots of data to obtain a parent distribution that approximates, or at least bounds, the wafer diffusion lot distribution. The resulting distribution (see figure 7) may favor one distribution type over the others, or it may not. However, lot-to-lot variability for some parts may be sufficiently large that such a procedure may not work. (See figures 8 and 9.)

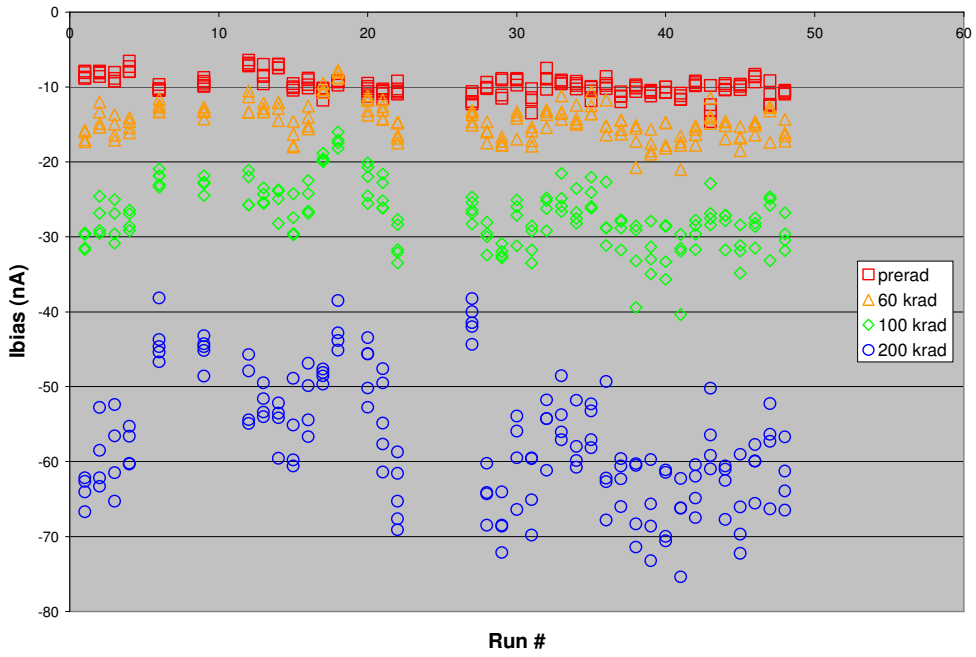

Figure 6 Data for 31 RLAT samples of the Linear Technologies RH1014 quad op amp exhibits little more spread lot-to-lot than within a lot. This means that data from multiple RLAT samples can be combined to constrain the intra-lot radiation response. Data supplied courtesy of J. Gorelick, Boeing Space Systems.

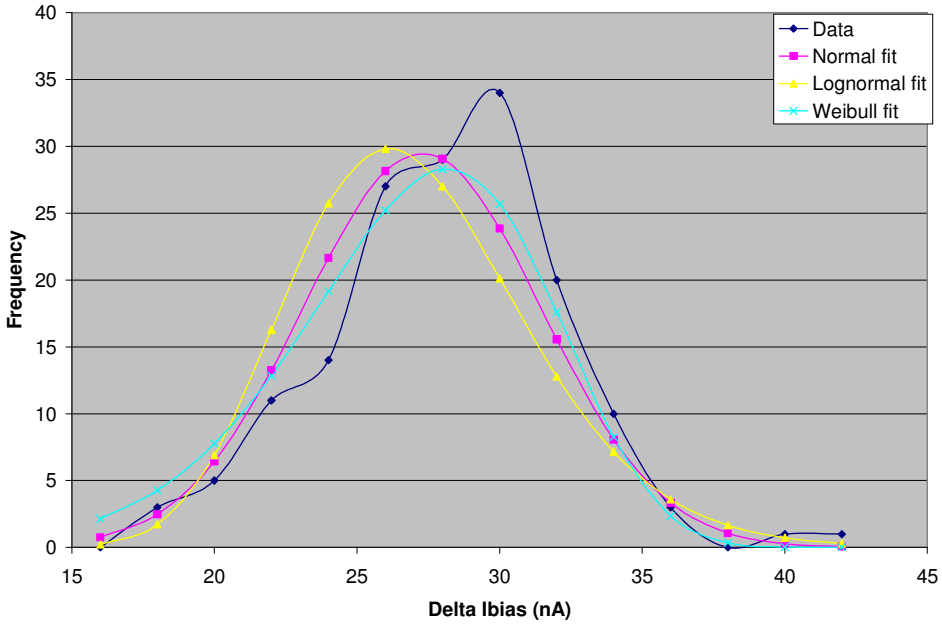

Figure 7 Change in Ibias for 154 pieces (31 RLAT samples) of RH1014 op amps after 100 krads(Si), along with Normal, Lognormal and Weibull fits to the data. Even for such a large dataset, no particular form is favored over any other.

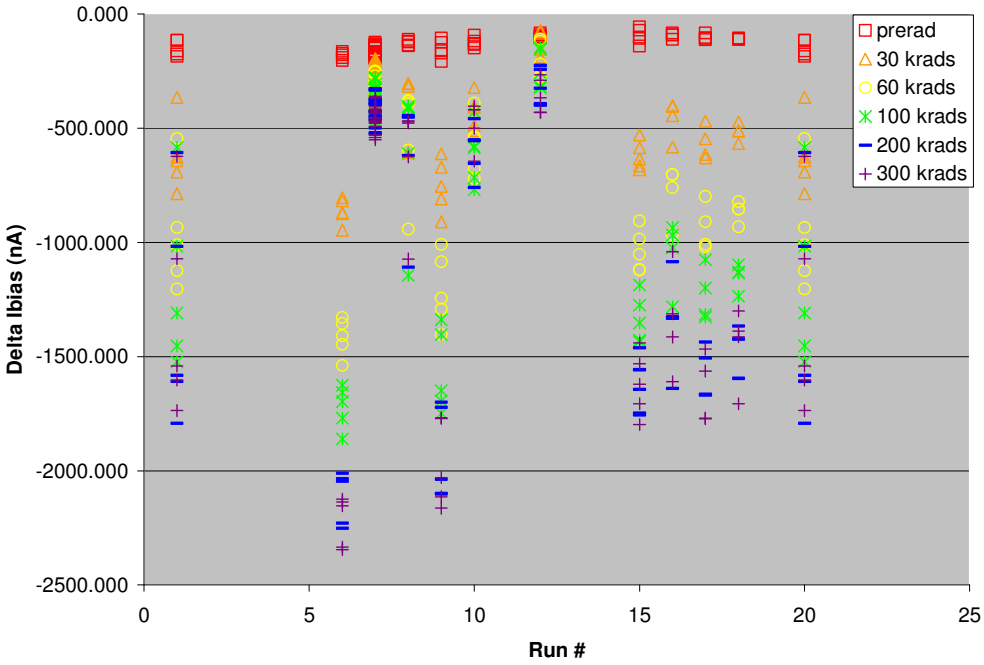

Figure 8 Data for 31 RLAT samples of the Analog Devices OP484 quad op amp exhibits considerably more spread lot-to-lot than within a lot. This makes it much more difficult to use archival data to constrain the intra-lot radiation response. More serious, the OP 484 exhibits bimodal response—perhaps even within a single wafer lot. Data supplied courtesy of J. Gorelick, Boeing Space Systems.

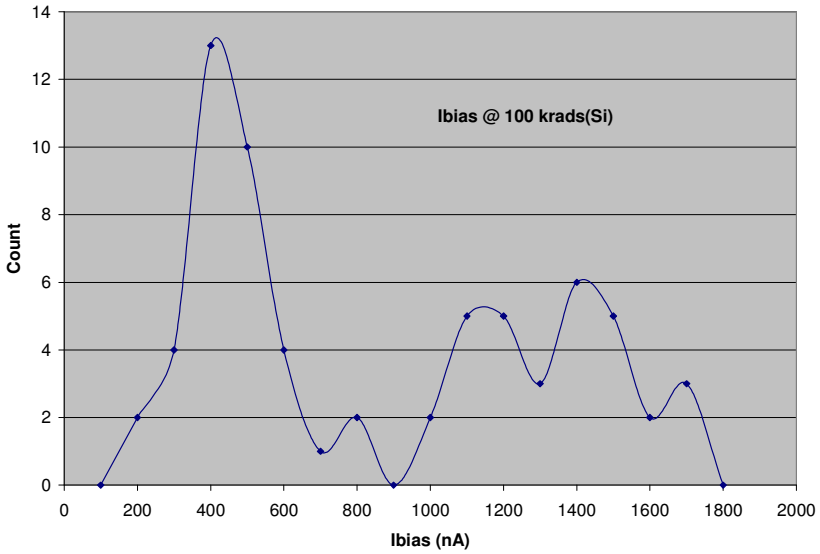

Figure 9 The OP484 op amp exhibits bimodal radiation response that worsens with increasing dose. For the most part, each RLAT sample correlated with only one mode. However, two lots had parts that fell into both modes.

Usually there will be no compelling reason to favor one distribution form over the others. In this case, since a small RLAT sample will not be sufficient to decide between the distributions, it is prudent to take the distribution that yields the worst case results. The spread of results over the three distribution types yields a measure of the systematic error that may result from the choice of distribution type. Unfortunately, such a comparison will not provide an accurate estimate of systematic error if the actual failure distribution varies grossly from the three distribution types assumed above. The most pernicious of such variations include distributions with very thick tails (e.g. Cauchy, Pareto, etc.) and bimodal (or multimodal) distributions. For the most part, distributions with very thick tails tend not to occur in failure studies, and can be eliminated from consideration, although they sometimes appear in answer to questions like, “Well, how bad can it be?” (thus, the use of the Cauchy distribution in Figure 2.) Multimodality, on the other hand, is known to occur in TID degradation. Usually, a single wafer diffusion lot will be associated with only a single mode. However, intra-lot bimodality does occur—two of the lots of OP484 used for figures8 and 9 had parts in both modes evident in figure 9. Multimodal distributions are particularly pernicious for several reasons. First, if one mode is much less probable than the other(s), a test with small sample size is unlikely to detect the less probable mode. This means that hardness assurance predictions for a large flight lot would likely not be based on a representative sample. Multimodality also renders increased RDM ineffective as an RHA strategy. Moreover, once a second mode is found, it is nearly impossible to preclude the existence of a third mode, and so on. As such, it is important to understand how bimodality can occur and how one can provide assurance against it. In what follows, we examine how archival data—if available and of sufficient quantity—may be used to increase confidence in the assumption of a single mode for the failure distribution. Where this is impossible, or where the analysis is not sufficiently conclusive, we examine aspects of radiation mechanisms and of device characteristics and fabrication, particularly as these affect susceptibility to bimodal radiation response.

4.0 Bounding the Radiation Response Distribution

Although systematic error estimates can be obtained by comparing results for analyses under assumptions of different distribution forms for the failure distribution, the existence of multiple modes in the failure distribution would invalidate these and all other analysis results. As such any analysis that increases the confidence in the single mode hypothesis, significantly increases confidence in RHA results. Here we consider two strategies: 1)Use of archival data, and 2)examination of the factors that may cause or contribute to bimodal radiation response in microelectronics.

4.1 Use of Archival Data

Although economical constraints on RHA prohibit testing very large samples, lot-to-lot variability often necessitates repeated testing of the same part types. As such, over time large statistics may accumulate for a given part type. Such datasets can serve several purposes. Since each RLAT sample provides an estimate of the failure distribution for its wafer diffusion lot, a dataset accumulated over time can indicate trends in radiation performance. Perhaps more interesting, it can help to constrain the parent failure distribution. If the process does not change significantly with time, then the parent failure distribution for a given lot of parts should be bounded by the lot-to-lot failure distribution. Sometimes—especially for radiation-hardened or military-grade parts—lot-to-lot variability may be comparable to intra-lot variability. Archival data may then reasonably approximate the intralot distribution. The archival data may contain a wealth of valuable information, in addition to the level of inter-lot and intra-lot variability. If there are sufficient samples available, one may even get an idea of the shape of the distribution, including not just its location and width, but higher moments and even its extreme value behavior. Independent of the form of the distribution favored by the archival data, such data can be used in conjunction with binomial statistics to place limits on the amount of bimodality in the parent distribution. For instance, the 158 entries for RH1014s shown in figures 6 and 7 are sufficient to limit the size of any second mode to 1% or less of all parts with 89% confidence.

Archival data can also be used to determine whether additional lot testing represents significant added value, and to determine the appropriate additional margin necessary if RLAT is to be waived. Treatment of such archival data sets can be either parametric (e.g. maximum likelihood fitting) or nonparametric (bootstrapping, randomization, jackknifing, etc.).[16] However, usually even under the best of circumstances, data are limited and to reach meaningful conclusions one must extrapolate beyond the support of the available dataset. In this case, the only choice is parametric analysis. Indeed, for most parts—especially commercial parts—there will be little or no archival data available. In this case, it is necessary to examine what is known of the part's technology, fabrication, and radiation effects mechanisms to determine whether bimodality is likely for the part.

4.2 Part Technology and Radiation Effects Mechanisms

As a last resort, in the absence of any other information eliminating pathological distributions from consideration, one is left with heuristic considerations related to device structures and fabrication and to radiation effect mechanisms as a measure of such risks. While a detailed discussion of such matters is beyond the scope of this study, the reader is directed to reference 17 for a discussion of semiconductor fabrication processes, to references 18-21 for discussions of TID mechanisms, reference 22 for displacement damage and to references 23 and 24 for discussions of destructive and nondestructive SEE mechanisms, respectively.

Unless the process by which a part is fabricated is tailored specifically to control variability of radiation response, then the risk of such variability is likely significant. However, the likelihood of significant variability is not equal for all radiation effects. Nor is the effect of the variability likely to be the same. As mentioned above, the fact that current SEE rate calculation routines are robust to small variability and have roughly a 2x systematic uncertainty besides means that part-to-part variability—and even bimodality—are unlikely to be noticed unless they are severe. In addition, the extent of variability is likely to be less if the structures and processes involved in the effect are the same as those involved in normal operations of the part.[25] For example, a SEU occurs in a CMOS device when charge collected as a result of the passage of an ion exceeds some critical value. Ignoring charge collected by diffusion for the moment, the device structures and processes involved in this process are very similar to those involved in normal device operation. The main difference is the source of the charge. As such, if there were too much

variability in SEU response, it might also affect the electrical performance of the part. Similar arguments might be advanced for SEFIs and SETs, since these also involve structures and processes typical of normal device operation. As might be expected, there are caveats associated with this optimistic picture. One of these is charge collection due to diffusion. Diffusion depends on the properties of the substrate (usually the same as those of the semiconductor wafer), which in many parts does not enter into the normal operation of the device. Thus, to the extent that diffusion contributes significantly to charge collection and to the extent that wafer properties vary within a wafer lot, one might expect to see a comparable degree of variability in the charge collected by diffusion. Having said this, great care is exercised in the growth of silicon wafers as well as in their characterization (with 32 or more characteristics per wafer being measured).[26] This makes variation in diffusion characteristics an unlikely source of bimodality. A second caveat derives from the fact that the smallest-sized features in current generation devices may now have only of order 100 dopant atoms, leading to significant statistical variation[27]. However, again these variations will tend to vary more or less smoothly, and the devices with the smallest feature sizes (e.g. access FETs in DRAMs) This coupled with the systematic uncertainties associated with current rate calculation methods makes it unlikely that variability—particularly bimodality—will affect the conclusions of SEU analyses.

The question of variability and bimodality for destructive SEE modes is somewhat more problematic, since these modes invariably involve processes not encountered during normal operation.

The errors on calculated rates for destructive SEE are usually even greater than those for nondestructive SEE rates. In part, this is because destructive SEE cross sections tend to be based on relatively few events, and so have large random errors. Because destructive SEE typically have higher onset LETs than nondestructive rates, the cross section curves also typically have fewer points to constrain the fit to the data. There is also the fact that the charge collection volumes for these destructive modes are invariably not rectangular parallelepiped (RPP), so the assumption of RPP volumes implicit in the rate calculation algorithms introduces systematic errors into the process. Thus, calculated rates for destructive SEE tend to err on the side of caution, and usually such conditions are rare enough on orbit that any errors tend to be obscured by Poisson fluctuations in any case. Nevertheless, there is reason for concern that part-to-part variability may become important as more accurate SEE rate calculation methods are introduced. Modes such as SEL and SEB involve parasitic structures that play little role in normal operation of the device. Modes such as SEGR depend on the differential mobilities of electrons and holes. However, while significant part-to-part and lot-to-lot variability have been seen, to date there is no evidence of bimodal behavior within a single wafer diffusion lot.

On the other hand, part-to-part variability and bimodality have a venerable history and may have serious consequences for TID degradation[18]—starting with so-called maverick parts and continuing through observations of bimodality associated with enhanced low dose rate sensitivity (ELDRS). The causes of bimodality in TID response are only partly understood. Krieg et al. discuss both the characteristics and RHA consequences of bimodality in LM111 comparators.[28] Ultimately, the multiple modes were attributed to differences in the thermal histories (so-called Pre-irradiation Elevated Thermal Stress—PETS) of otherwise identical parts, with these differences shifting the balance between competing degradation mechanisms and giving rise to markedly different ELDRS and TID responses. [29] In some ways, this is not surprising, since TID response depends on details of Si/SiO2 interface, as well as trap densities in the bulk oxide—both of which can be affected by thermal stresses. Moreover, while large changes in interface roughness can affect charge transport properties (particularly mobility) the subtle changes that give rise to different TID responses are not generally of great importance for how well an

oxide fulfills its functions as an insulator. For state-of-the-art semiconductor devices the use of rapid thermal annealing (RTA) may also be a source of variability, since it is necessarily a nonequilibrium process. (RTA is used to activate dopants and repair damage due to ion bombardment while preserving sharply graded dopant profiles.) As such, it is not hard to understand how parts might exhibit significant variability even within a wafer lot. Introduction of bimodality is somewhat harder to understand unless some difference in processing (e.g. PETS or differential heating during RTA) gives rise to differential response. Bimodality does occur, however, in a small minority of parts. In some cases, it manifests quantitatively—with parts in one mode exhibiting more degradation for a given dose. In other cases, different modes may exhibit qualitatively different response. For instance, the Motorola ?????? op amp from Motorola (now obsolete) typically exhibited worse degradation when unbiased than when biased. However, one part in a lot of 3 unbiased parts exhibited degradation very similar to that of biased parts.[30]. Indeed, this last example illustrates one of the problems inherent in discussions of bimodal radiation response: with a single small sample, it may be impossible to distinguish between bimodal response and a very broad single mode.

Qualitative differences have also been seen in the displacement damage response of linear bipolar and electro-optical devices [31],[32], although to date there is no evidence of bimodal response within a wafer diffusion lot. As long as the wafers used in a single diffusion lot do not vary significantly, and as long as processing steps that affect lattice quality are similar, bimodality is probably more unlikely for displacement damage than for TID. Similarly, although SEL susceptibility has been seen to vary significantly from lot to lot for some DRAMs, bimodal response has not yet been seen within a single wafer lot.

Although to date bimodality has only been seen in the TID response of only a small minority of parts, it poses significant challenges for any application employing large flight lots. Small RLAT samples have very little chance of catching low-probability modes, and even if they do contain samples from all modes, the analyst will generally be unable to distinguish whether the “outlier” is due to a second mode or to an excessively broad distribution. On the other hand, large flight lots are very likely to contain parts from all modes. The only comforting news is that so far the parts that have exhibited significant bimodality (bipolar junction transistors, op amps, comparators, etc.) are not the same parts that are likely to be flown in very large flight lots. Indeed, unless an application had very stringent requirements, it could fulfill its demand for large numbers of such parts with parts known to have a unimodal response. Still, a better understanding of how bimodal response arises in TID degradation would significantly improve RHA for large flight. Nevertheless, it is not unreasonable to continue assuming unimodal response—especially since there is no economically viable alternative to doing otherwise.

5.0 Random and Systematic Errors for Large Flight Lots

As with any other RHA methodology, RHA for large flight lots begins with the performance requirements for the application. Among the most important of these is what constitutes a failure *for the application*. Considerations that are important in such specifications include not just application performance and parametric requirements, but also the level of redundancy in the application. In some cases, failure of a single part out of N parts in the flight lot may constitute a failure of the application. In other applications, there may be significant redundancy (N for M redundancy), so that several (N-M+1) parts must fail before the application fails. Radiation requirements should take this redundancy into account. For instance, if the failure of a single device represents a system failure, then the probability of interest is that of the first failure of N parts occurring for a stress S. If the probability of failure for any of the N parts in the flight lot is $p(S)$, and the cumulative probability is $P(S)$, it is easy to show that the probability of interest is just

$$p(1st\ failure\ @\ S) = \binom{N}{1} p(S) \times (1 - P(S))^{N-1} \quad (8)$$

Note that when the cumulative probability of failure is small, this reduces to *N∗ p(str)* as expected. The cumulative probability that at least one failure has occurred for stress *S* is just 1 minus the probability that no failures have occurred out of the *N* parts:

$$P(>0\ failures) = 1 - P(0\ failures) = 1 - (1 - P(S))^{N} \quad (9)$$

If the system has *N* for *M* redundancy, system failure will occur after the *(N-M+1)th* failure, and the probability of seeing this failure at stress *S* is

$$p((N - M + 1)th\ failure\ @\ S) = \binom{N-M}{1}\binom{N}{N-M} (P(S))^{M}\ p(S) \times (1 - P(S))^{N-M-1} \quad (10)$$

(see figure 10) and the cumulative probability for system failure (i.e for seeing at least *N-M+1* failures) is

$$P(\geq N - M\ \ failures\ @\ S) = 1 - \sum_{i=0}^{N-M} P(i\ \ failures) = 1 - \sum_{i=0}^{N-M} \binom{N}{i} (P(S)^{i} (1 - P(S))^{N-i} \quad (11)$$

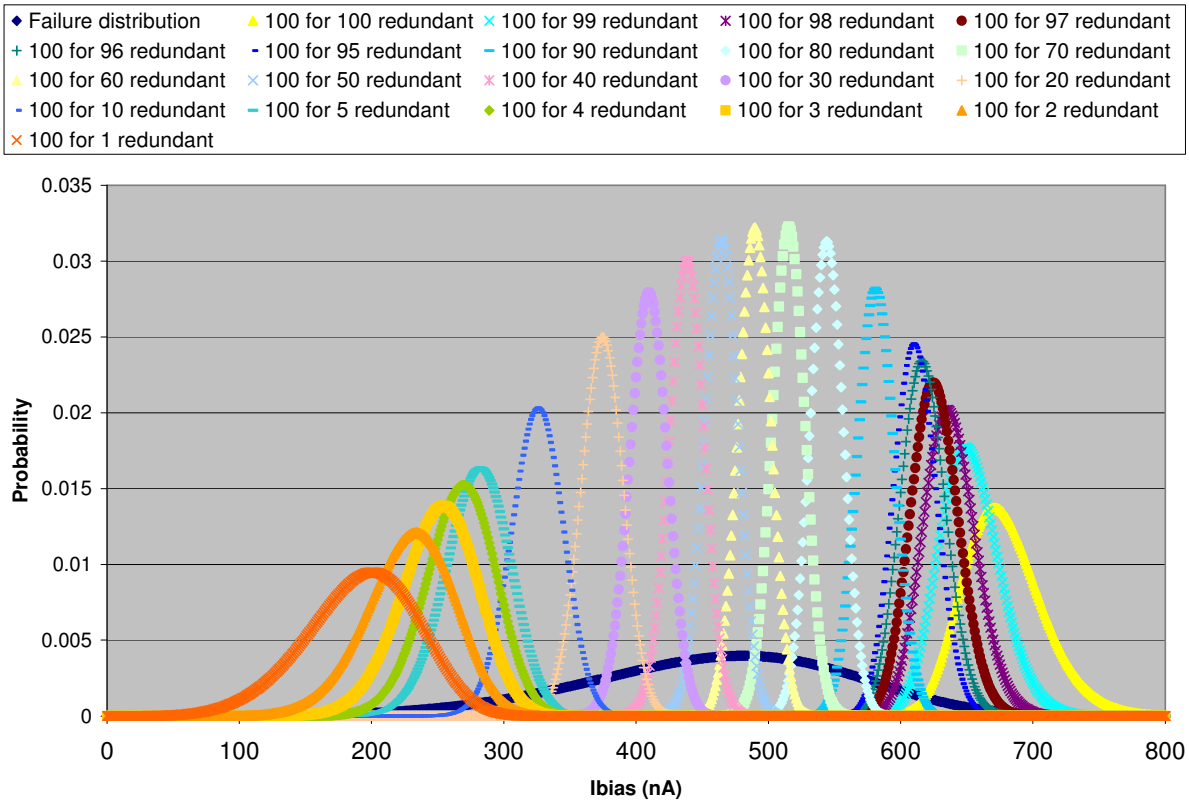

Figure 10 The effect of redundancy on the system failure rate depends on the underlying failure distribution (solid blue curve), the number of parts in the flight lot (100 for this case), and the level of redundancy. As the level of redundancy approaches a ratio of ~2:1, the peak of the failure approaches the median of the underlying failure distribution.

It should be pointed out that for large flight lots and for high levels of redundancy, the combinatorics involved in these formulas can become intractable, exceeding the accuracy possible with 32-bit processors. In such circumstances, it may be feasible to use Stirling's approximation for the factorial

$$N! \approx \sqrt{\frac{\pi}{3}(6N+1)}N^N e^{-N} \qquad (12)$$

Note that in this approximation, the product $N!$ is replaced by much more compact expression. This means that large and small factors can be combined to achieve a result as close as possible to 1 prior to being raised to a large power. By judiciously combining factors it is possible to achieve reasonable approximations for flight lots as large as 100000—far beyond any current levels of complexity.

In setting radiation requirements, it is important to ensure that the probability considered takes into account the appropriate redundancy for the system. In what follows, we assume no redundancy, since this constitutes a worst case. The results can be extended to redundant systems by substituting the appropriate probability expression for the zero-redundancy probability expression above (10-11). In addition, one must specify an acceptable probability of failure P(F) for the application based on mission requirements and a confidence level CL at which it must be demonstrated that the failure probability is less than the required level. With this information, one is ready to plan the test program, beginning by determining the sample size to be tested.

5.1 Determination of Sample Size and other Guidelines for Testing

If the failure distribution is Normal, then it is fairly straightforward to determine of the appropriate sample size using the Student's t distribution, as outlined in section **3.0**. For distributions other than the Normal, no equivalent to the Student's t distribution exists—or rather, an infinite number of such distributions exist, since the small sample convergence properties of the distribution depend on the values of the distribution parameters, themselves. As mentioned briefly above, for given values of the fit parameters, one can estimate the equivalent to the Student's t distribution numerically. While this may seem to be of little help given that one may not know the distribution parameters a priori, one can generally restrict the parameters to those corresponding to distributions with standard deviations less than 30% of the mean. Broader distributions will generally be problematic for hardness assurance in any case, since they imply significant variance in response, and will be likely to yield poor estimates of mean response. Moreover, for very narrow distributions ($\sigma/\mu<0.05$), the conclusions of the analysis tend to be insensitive to the assumed distribution unless the flight lot is extremely large (>1000 units). In any case, a narrow distribution in radiation response is a problem radiation analysts experience all too seldom. Thus, the range of parameters that must be considered is limited, and the behaviors of the distributions may be parameterized over this range with reasonable accuracy. (See figure 11, for example.)

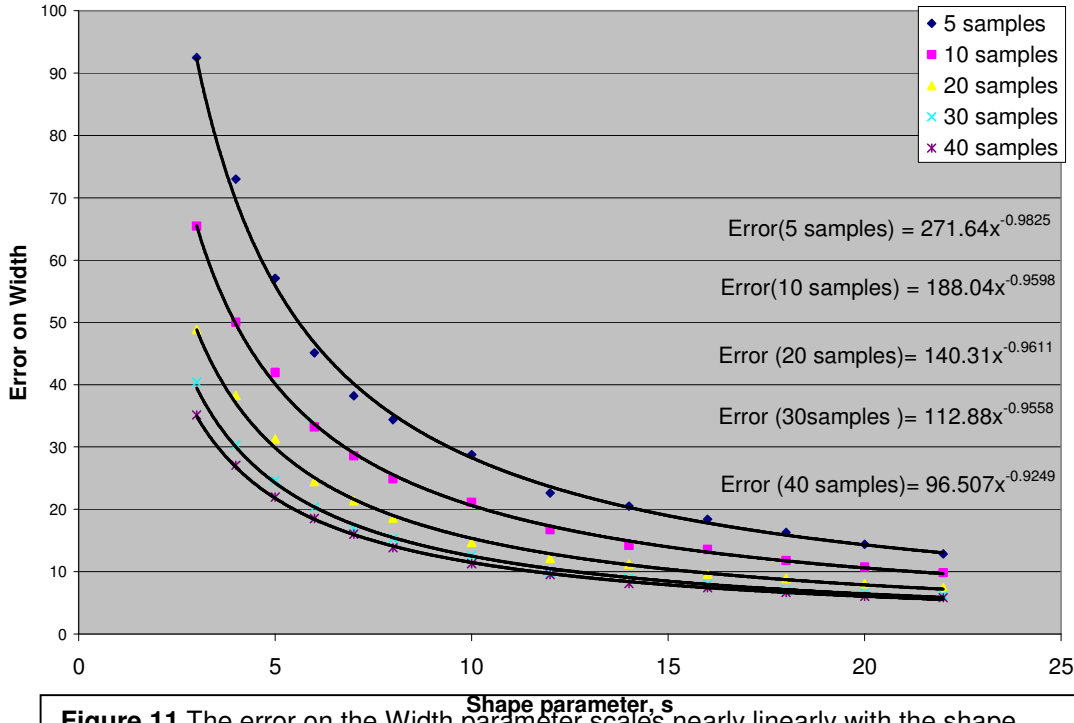

Figure 11 The error on the Width parameter scales nearly linearly with the shape parameter, s, and as the inverse square root of the sample size (as indicated by the proportionality constants in the power laws).

The task is further simplified by the fact that the ratio μ/σ is independent of the Weibull width parameter and of the logarithmic mean μ_{ln} for the lognormal. As a general rule, the convergence of the parameters that most influence the location of the distribution (mean, μ, in the Normal, μ_{ln} for the lognormal and Weibull width parameter, W) tend to converge as the inverse square root of the number of samples. Parameters that most affect the higher moments of the distribution (e.g. the Weibull shape, σ_{ln} and even the standard deviation of the Normal) tend to have a somewhat more complicated and often slower convergence with sample size. For example, if the distribution is particularly narrow, increased sample size will aid convergence of estimators for the Weibull shape or σ_{ln} only to a limited degree.

The asymmetry of the lognormal and Weibull also mean that upper and lower confidence limits will also likely by asymmetric and will probably converge differently as sample size increases.

In addition to the convergence properties of the assumed form(s) for the distributions, required sample size is determined by several factors, including the criticality of the application (which determines the acceptable failure level and confidence level) and the tolerance of the application for degradation (which determines the tolerance for error on estimators for failure distribution parameters). In general, the method outlined in Section 3.0 represents a good first guess for the required sample size. One can then look at the convergence properties of the lognormal and Weibull distributions in for the anticipated (or worst case) range of parameters, and if it is significantly worse than for the Normal, consider augmenting the sample.

Although we have been concerned here mainly with the analysis of data, different types of data are more or less valuable in constraining a failure or radiation response distribution. Clearly, for a failure distribution, the radiation increments should not be too far apart, while for a radiation response distribution it is crucial that the stresses applied to each part be equal. However, the method by which the data are to be analyzed also determines what data are of most value. For instance, if a

failure distribution is to be fit to a distribution form using the Maximum Likelihood method, then actual observed failure levels (as opposed to intervals or suspensions) provide greater constraints for the fit. This difference will be less important if data are to be fit by some other method (e.g. Least Squares, Method of Moments).
Once the test is performed on the required sample size, the data from the test can be analyzed and combined with whatever other information is available.

5.2 Data Analysis

While considerable effort usually goes into the design and execution of a radiation trial, data analysis is often cursory. In part, this is because the implications of the trial for the application are often unambiguous. However, this cannot be said of the effort to extrapolate the results for a small test sample to the performance of a large flight lot. In this case, extreme caution must be taken to ensure that the conclusions drawn are as rigorous as possible, since large flight lots will magnify the effects of any errors in analysis.
In general, the preferred fitting method is Maximum Likelihood Estimation (MLE), since this yields confidence intervals as well as estimates for parameters. Moreover, the MLE method lends itself to numerical estimation, and can even be implemented in Microsoft Excel™. In what follows, we will use the same archival data for the RH1014 and OP484 that were shown in figures 6 and 8.

In many cases, an analytical, closed expression can be found for a distribution's ML estimators. However, it is often easier to solve for them numerically. In addition, this has the advantage of revealing the behavior of the likelihood as one moves away from the best estimate. In turn, this allows estimation of the confidence contours for the estimators.
Because ML estimators are generally normally distributed about the actual distribution parameter values, the likelihood equation tends to vary as the χ^2 function for a given confidence level 1-α and number of parameters allowed to vary ν:

$$-\Delta \ln(L(\Theta_1, \Theta_2, ...) \approx 0.5\chi^2_{\alpha;\nu} \qquad (12)$$

Figure 12 shows the likelihood for a Weibull fit to changes in leakage current for the RH1014 after the 200 krad(Si) dose step. The plot indicates not only the best fit (magenta), but also the 90% (inner edge of yellow contour), the 95%CL (outer edge of yellow contour) and the 99% and 99.9% CLs (inner and outer edge of the green contour, respectively).

	0	1	2	3	4	5	6	7	8	9	10	11	12	13
width\shape	6.06	6.36	6.66	6.96	7.26	7.56	7.86	8.16	8.46	8.76	9.06	9.36	9.66	9.96
58.0125	-392.12	-390.52	-389.55	-389.21	-389.48	-390.34	-391.79	-393.80	-396.37	-399.51	-403.20	-407.45	-412.25	-417.61
58.3125	-391.16	-389.36	-388.18	-387.61	-387.62	-388.19	-389.32	-390.98	-393.19	-395.92	-399.17	-402.94	-407.23	-412.04
58.6125	-390.37	-388.40	-387.04	-386.26	-386.04	-386.35	-387.20	-388.56	-390.43	-392.79	-395.65	-399.00	-402.83	-407.15
58.9125	-389.76	-387.65	-386.12	-385.16	-384.73	-384.82	-385.42	-386.51	-388.08	-390.12	-392.63	-395.59	-399.02	-402.89
59.2125	-389.32	-387.08	-385.41	-384.29	-383.68	-383.58	-383.96	-384.81	-386.11	-387.87	-390.06	-392.69	-395.75	-399.23
59.5125	-389.03	-386.69	-384.91	-383.65	-382.89	-382.61	-382.80	-383.44	-384.51	-386.02	-387.94	-390.27	-393.00	-396.14
59.8125	-388.90	-386.48	-384.59	-383.21	-382.32	-381.90	-381.93	-382.38	-383.26	-384.54	-386.22	-388.29	-390.74	-393.57
60.1125	-388.92	-386.42	-384.46	-382.99	-381.99	-381.44	-381.33	-381.63	-382.33	-383.42	-384.90	-386.74	-388.95	-391.51
60.4125	-389.08	-386.53	-384.50	-382.95	-381.86	-381.21	-380.98	-381.16	-381.71	-382.65	-383.94	-385.59	-387.58	-389.91
60.7125	-389.37	-386.78	-384.70	-383.10	-381.94	-381.21	-380.89	-380.95	-381.39	-382.19	-383.33	-384.82	-386.63	-388.76
61.0125	-389.79	-387.18	-385.07	-383.42	-382.21	-381.42	-381.02	-381.00	-381.34	-382.03	-383.05	-384.40	-386.06	-388.02
61.3125	-390.33	-387.71	-385.58	-383.91	-382.67	-381.83	-381.38	-381.30	-381.56	-382.16	-383.08	-384.32	-385.85	-387.68
61.6125	-390.99	-388.37	-386.24	-384.56	-383.30	-382.44	-381.95	-381.82	-382.03	-382.56	-383.41	-384.56	-385.99	-387.71
61.9125	-391.76	-389.16	-387.04	-385.36	-384.10	-383.22	-382.72	-382.56	-382.73	-383.22	-384.01	-385.09	-386.45	-388.08
62.2125	-392.63	-390.07	-387.97	-386.31	-385.05	-384.18	-383.67	-383.50	-383.66	-384.12	-384.87	-385.91	-387.22	-388.79
62.5125	-393.61	-391.08	-389.02	-387.39	-386.16	-385.31	-384.81	-384.64	-384.79	-385.25	-385.99	-387.00	-388.27	-389.80
62.8125	-394.69	-392.21	-390.19	-388.60	-387.41	-386.59	-386.11	-385.97	-386.13	-386.59	-387.33	-388.34	-389.60	-391.10
63.1125	-395.86	-393.44	-391.48	-389.94	-388.79	-388.01	-387.58	-387.47	-387.66	-388.14	-388.90	-389.91	-391.18	-392.68
63.4125	-397.12	-394.77	-392.87	-391.40	-390.31	-389.58	-389.20	-389.13	-389.37	-389.88	-390.67	-391.71	-393.00	-394.51

Figure 12 As with the fit for SEE data, the MLE fit Ibias data for the RH1014 after 200 krad(Si) yields best estimates for the fit parameters (Weibull scale and shape in this case). The inner edge of the yellow contour corresponds to the 90% CL, and the outer to 95% CL, and the inner and outer edges of the green contour correspond to the 99% and 99.9% CLs, respectively.

The same data can be fit simultaneously to a lognormal and Normal distribution. Although the data indicate that the best fit to the Weibull is for W=60.71, s=7.86, this fit is susceptible to statistical fluctuations. As such, it is prudent to take the fit parameters that yield the worst-case distribution (largest breadth and mean) and are consistent with the confidence level prescribed by the requirements. For the Weibull fit shown in figure 12, the 95% CL corresponds to a width parameter ~62.4 and shape parameter ~6.7 (as opposed to best fit values W=60.7125, s=7.86). Likewise, the lognormal likelihood yields a best fit for μln=4.04, σln=.1528 and 95% WC fit at μln~4.06, σln~.177, and the Normal likelihood yields a best fit for μ=57.2, σ=8.4 and 95% WC at μ~58.5, σ~9.8.

As an example of how such fit information might be used, consider the assignment of design-to values for increased bias current. We assume there is no redundancy and that any failure represents a failure at the system level. Fits similar to that shown in figure 12 for the 200 krad(Si) radiation step are conducted for the 60 krad and 100 krad data. We assume that the flight lot contains 20 parts, and that the acceptable failure probability is 0.3% at the 95% CL, and plug the resulting failure distributions into equation 12. Figure 13 shows the results, including the best-fit and 95% WC Weibull and Lognormal fits and the corresponding system-level failure distributions for a 20-part flight lot. (The normal fit is suppressed for simplicity and because the results are intermediate between the Weibull and lognormal.)

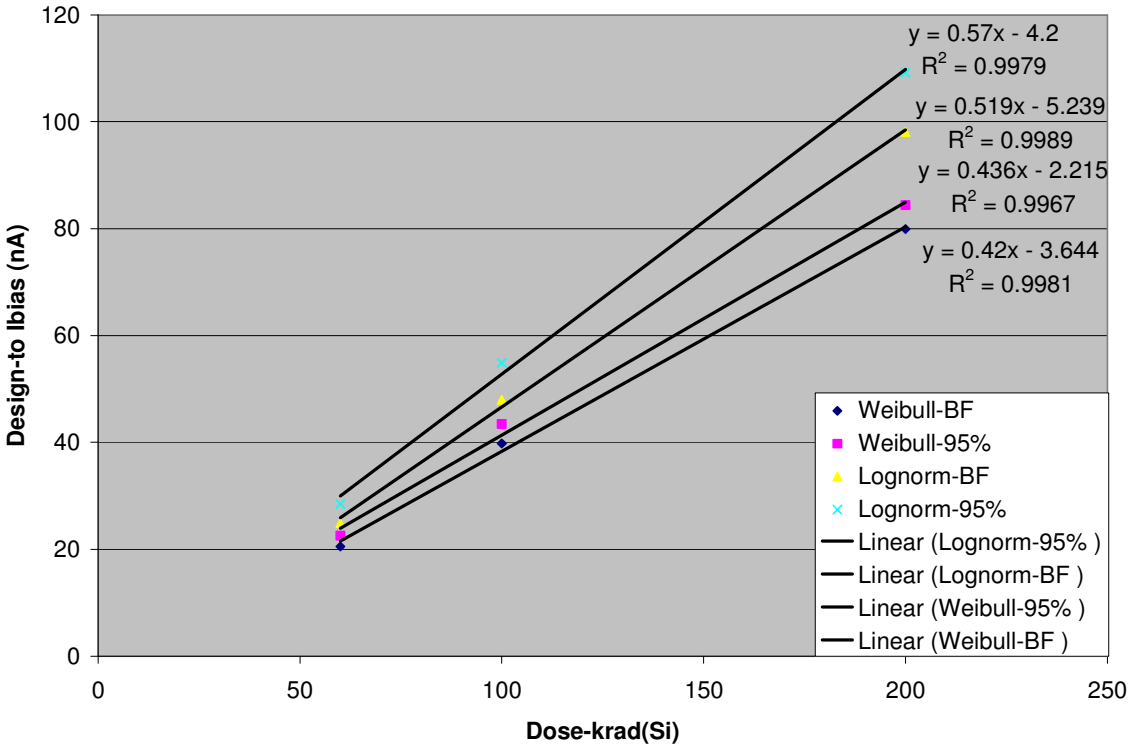

Figure 13 Design-to values for RH1014 data analyzed in this example give bias currents that would only be exceeded 0.3% of the time by any of the parts in a randomly selected flight lot of 20 parts.

The condition defining the design-to bias currents is that there is less than 0.3% probability of drawing a part from the failure distribution that exceeds the design-to values. Figure 13 shows the design-to values for the best-fit and 95% WC curves. The choice of distribution affects results at about the 20-25% level, while the choice of the 95% CL over the best-fit values has a 5-10% effect. The linear trend of increased bias current with increasing dose makes it possible to implement some simplifications to the above procedures. For instance, the fact that the 95% WC

design-to values are about 2x the mean values for the data on which they were based means that if one imposes an additional 2x margin on top of the standard 2x RDM (4x total), then the available archival data may provide adequate assurance (even without lot testing) that the parts will meet their design goals. This presumes, of course, that the archival data are representative of the flight lot under consideration and that the test conditions for the data are appropriate given the application conditions—assumptions that must be verified by independent means. This procedure provides rigorous justification for waiving RLAT for applications with adequate margin—in contrast to the usual one-size-fits-all approach that assigns a firm RDM requirement for all parts and ignores detailed implications of archival data.

Unfortunately, not all parts can be treated in this fashion. As discussed previously, the bimodality of the OP484 data in figures 8 and 9 precludes use of added RDM as a hardness assurance strategy. If we ignore the risks and precede anyway by fitting the upper mode to Weibull, Normal and Lognormal forms, we still run into problems. First, dividing the data increases statistical uncertainties. In addition, Figure 14 shows the resulting plot of design-to Ibias as a function of dose. Clearly, it is nonlinear, making it much more difficult to correlate increased RDM with increasing hardness assurance. While this is not an insurmountable barrier (i.e. one can derate based on the worst case design-to current or assume piecewise, worst-case linear fits between dose steps), it significantly complicates the task of analysis.

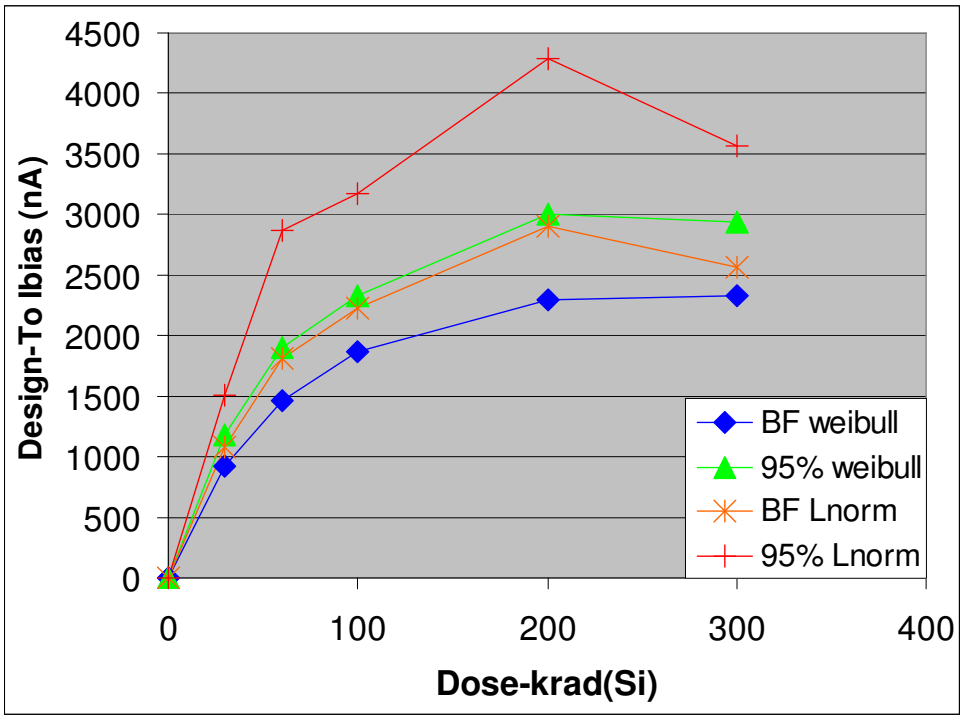

Figure 14 Saturation of Ibias with dose means that the Ibias vs. dose is nonlinear, precluding a simple analysis to determine an equivalent RDM for any given confidence interval. In this case, the best-fit and 95% WC values differ significantly, indicating that the fit for the data by either form is marginal.

Unfortunately, the design-to values reflect the difficulty of fitting this data—both the choice of 95% CL and distribution type can affect the design-to values at the 45-50% level. While this may seem rather punitive, it reflects the uncertainty in the underlying data. More significant, for the highest doses tested, the OP484's distribution of leakage current shows what may be the beginnings of still more modes in the distribution (see figure 15). Overall, use of archival data for a part with behavior like

that seen in the OP484 would be inadvisable. However, the techniques outlined here can also be used with RLAT data specific to the flight lot.

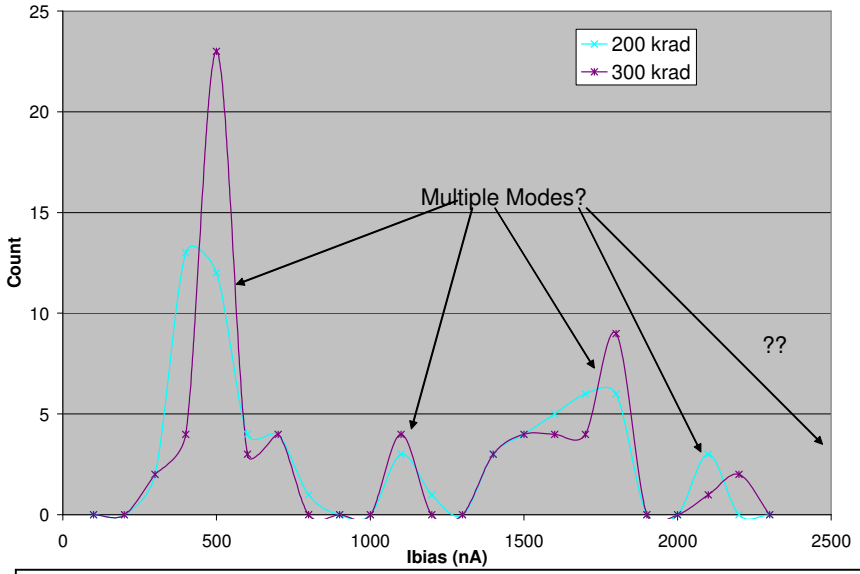

Figure 15 For the 200 and 300 krad(Si) dose steps, the OP484 begins to show indications that the high-mode seen at lower dose steps is fragmenting even further. Data such as this means that little additional assurance can be gleaned by examining archival data—meaning RHA fo this part depends critically on RLAT.

Although the flight lot size assumed here is moderate, it is clear that the distribution assumed and the confidence level chosen for the analysis do affect the design-to values (~20% and 5-10%, respectively, in the example above). These effects increase with flight-lot size.

In most applications sufficient data will not be available to distinguish which distribution form is the "right" one. In reality, the actual failure distribution may not conform to any of the forms considered here. In such cases, it is prudent to assume the form that yields worst case results, and to take the difference between the results for different forms as a measure of the systematic uncertainty introduced by assuming a particular form for the distribution. The resulting estimate of systematic error should be a reasonable approximation as long as the actual failure distribution is relatively well behaved (e.g. not multimodal or thick-tailed, etc.).

5.3 The Effect of Redundancy

Introducing redundancy into the system can significantly increase its reliability—or, equivalently, ease the design-to levels required to ensure that reliability. In the RH1014 example discussed above, the addition of a single redundant part reduces the design-to currents by up to 20%, depending on the distribution assumed (larger for lognormal than for Weibull) and whether one uses the best-fit distribution or the 95% WC distribution. The effect of adding additional redundancy diminishes, but the difference between the assumed distributions and for different confidence levels also decreases. (See figures 16 and 10.)

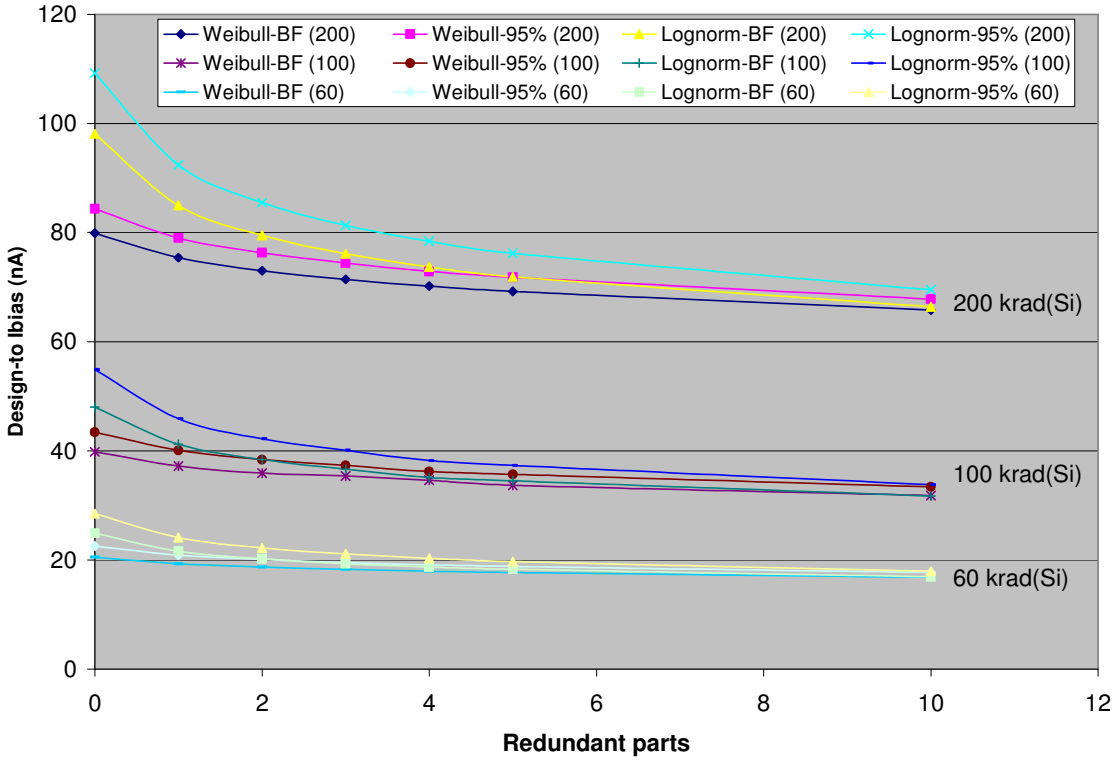

Figure 16 Adding redundancy increases the reliability of an application—or equivalently allows achievement of a comparable lifetime with a relaxed design-to value.

That is, increasing redundancy decreases the dependence of system performance on the exact form assumed for the failure distribution. Even (or rather especially) for pathological thick-tailed and bimodal distributions, the improvement is significant. This rather intuitive result is also easy to understand mathematically. Additional redundancy means that the system performance is driven less by the worst-performing part (the extremes of the distribution) and more by its central behavior—it's mean, mode, and median. For sufficiently large flight lots and redundancy levels, this behavior will appear Normal.

The central dependence of redundant system performance makes such systems inherently more robust (at least in principle) than nonredundant systems—independent of the actual form of the failure distribution. Unfortunately, redundant systems are also inherently much more complicated than nonredundant ones. They also entail a weight penalty that may make them unacceptable in some space flight applications. Implementation of such systems presumes the ability to isolate failed components and to swap in replacements—and the switching elements required for these tasks can also fail. Careful design and analysis are required to ensure that the redundancy scheme increases the overall reliability of the system significantly.

Although the above discussion of part-to-part variability has mainly been concerned with TID degradation and other degradation mechanisms, the considerations of redundancy can also be applied to both destructive and nondestructive SEE. However, there are a few differences. In the case of destructive SEE, the probability of failure is effectively zero for an unbiased part, so the use of cold spares can significantly increase reliability and life of the application. If for some reason the use of cold spares is not possible, redundancy may not significantly increase reliability, since the probability of observing an event increases more or less linearly with the number of parts being flown (and biased). For nondestructive SEE, the use of triplicate-and-vote and similar schemes are examples of the redundancy strategies:

Since SEE are generally rare, the voting scheme effectively removes the extremes (e.g. SEUs) of part behavior. However, the same caveats about implementation apply, since now the voting circuitry represents a common vulnerable node.

6.0 Conclusions and Future Work

The goal of the preceding analysis has been to extend the usual methods of hardness assurance to systems flying very large flight lots of nearly identical parts or requiring exceptionally high reliability. In such systems, either the sheer number of parts being flown or the consequences of a failure can cause the risk analysis to be dominated by rare error and failure mechanisms that are difficult to detect with traditional RHA methodologies. An analysis of the main sources of error for RHA reveals that at least for the present, the predominant errors break down as follows: Poisson errors dominate SEE analyses and part-to-part variability is of greatest concern for degradation mechanisms—particularly TID degradation. (We note, however, that as SEE rate calculation algorithms improve in accuracy and as microelectronic devices continue to shrink, part-to-part variability will likely become more important for SEE, and TID and displacement damage will likely begin to exhibit a single-event character in addition to their current gradual nature.)

Treatment of the Poisson errors in SEE elucidates three main results. First, the possibility that a null result may represent a downward fluctuation, rather than an indication of immunity introduces the concept of an upper limit for the SEE cross section at a given confidence level. A figure of merit type upper limit for the rate can then also be estimated. Inversion of this idea introduces the concept of a minimum flux necessary to probe all potentially vulnerable structures on a part. While such a concept is simple in principle, for very small feature sizes, the complicated radial and temporal dependence of the charge density of the ion track must be considered. However, for the current generation of parts, ion fluences of 1E7-1E8 cm^{-2} should be adequate to establish immunity to an effect at the 95% CL.

Although it is more difficult to establish a confidence level for a positive result for an SEE trial, such confidence levels are desirable, since they make possible quantitative comparisons of SEE rates calculated from different data sets and by different analysts. They also facilitate the use of SEE rates in system reliability analyses. Although it is possible to use likelihood ratios to extrapolate from a given σ vs. LET curve to a bounding, worst-case curve consistent with a given confidence lever, the preferred method for estimating SEE rate confidence levels is use of maximum likelihood estimation to determine worst-case fit parameters consistent with the desired confidence level.

The third result derived for single-event effects concerned the additional complications posed by the low cross sections of proton-nucleus spallation reactions responsible for proton-induced SEE. The fact that perhaps one in 10^4 to 10^6 protons will produce recoil products of sufficient range and LET to produce an upset means that probing a device with protons to the same level as suggested above for heavy ions would require on the order of 10^{13} protons/cm^2—enough to deposit on the order of a Mrad(Si) of dose. While it is unlikely that any single part would see such a proton fluence in a single mission, very large flight lots increase the effective vulnerable area and so the probability of the effect. Thus, most parts would fail due to TID exposure before they exhibited the susceptibility to protons, but on orbit, the sheer volume of parts being flown makes it much more likely that such errors would be seen.

We then turned our attention to the random and systematic errors that may be introduced by attempts to deal with sampling errors in RHA. Because TID degradation is the place where part-to-part variability is most likely to affect system performance, we have concentrated on this threat in our discussion. The essential issue is that for very large flight lots or applications requiring very high reliability, one must consider not just the central behavior of the failure distribution, but also the behavior of the distribution tails. This makes it impossible to qualify a part for the system using a small sample lot without making assumptions about the part's failure distribution. Unfortunately, to the extent that the actual distribution differs from that assumed, such assumptions introduce systematic errors that may

be difficult to quantify. An examination of the sampling issues involved with RLAT and on-orbit performance prediction identifies the most serious pathologies of distributions to be bimodality and thick tails. (Relatively small sample sizes are often adequate to provide estimates of mean behavior and even the width of the failure distribution.) We have sought to demonstrate how one might allay concerns about such pathologies using archival data—or where this is unavailable or inadequate, available knowledge of the part's fabrication, device structures, materials and radiation effects mechanisms. These latter considerations identify TID degradation as the radiation effect most likely to be affected by serious part-to-part variability, and in rare cases, bimodality. Other radiation effects are expected to be less susceptible to variability within a wafer diffusion lot, because the structures and materials that give rise to these effects tend to be less variable than the oxides and especially the Si/SiO2 interface responsible for TID degradation.

If the effort to allay concerns about pathological failure distributions is successful, we then recommend analyzing the qualifying data for the part under the assumption of several different forms for the failure/radiation performance distribution. For our part, we have chosen to use three different distributions: 1)the Normal distribution, which is the easiest distribution to analyze and represents the limiting case for the central behavior of most well-behaved distributions; 2) the lognormal distribution, which is physically justifiable in some cases and can assume a wide range of morphologies; and 3) the Weibull distribution, which, like the lognormal is physically justifiable and assumes a range of forms, depending on its parameters. In addition, the three forms assumed here were chosen because they vary significantly in terms of the properties of their higher moments (skewness, kurtosis, etc.), allowing us to probe the sensitivity of our conclusions to variations in these properties. The range of predictions for the three forms assumed above represents a measure of the uncertainty introduced by assuming a particular form for the distribution—and if there is no compelling reason to favor one distribution over the others, prudence dictates that the worst-case distribution be chosen. In addition, the use of maximum likelihood estimation to derive not just estimators for distribution fit parameters, but also confidence intervals for them is also discussed. An example of how to use the above considerations in deriving design-to values—and in some cases equivalent radiation design margins—is discussed. Finally, the effects of adding redundancy to the system are discussed, and the rather intuitive conclusion is reached that redundancy improves system reliability by making the system performance depend more on the mean of the distribution than upon its extremes.

The method outlined here allows for conservative inference of the likely behavior of large flight lots from small to moderate-sized RLAT or other qualification samples. The advantages of the method are that it estimates random (both Poisson and Binomial sampling) and systematic errors. In particular, it outlines a method for estimating the potential errors introduced by assuming a particular form for the distributions being estimated, providing a degree of robustness against minor variations from the assumed form. Unfortunately, neither this method, nor any other method based on inference from small sample sizes, can provide robustness in the face of pathological failure distributions—particularly against those that are multi-modal. While the methods and arguments presented provide some reassurance that such pathological radiation response is likely to remain rare, it is clear that elucidation of the causes of bimodality in radiation response will significantly increase confidence in RHA analyses. This is particularly true for purely commercial components, where there is likely to be little hope for lot traceability.

In addition to a better understanding of bimodality, we believe that the current work can be extended in several ways. One obvious but difficult undertaking is to extend the analysis of part-to-part variability to SEE and to determine whether such variability may soon affect the conclusions of SEE analyses. A possible application for such an analysis is the long-term reliability of solid-state recorders using large quantities of commercial memories without lot traceability. Likewise, investigation of the potential Poisson nature of damage mechanisms in deep submicron devices represents another possible avenue for extension. Another possible study involves investigating of whether the techniques discussed here might be incorporated as part of a more general quasi-Bayesian methodology for radiation risk analysis. Some preliminary work has been done along these lines for the James Webb

Space Telescope. A Bayesian methodology is attractive for radiation analysis because it allows risk assessment to be based on all available information, while allowing for modification as new information (e.g. RLAT data) becomes available. Such a paradigm could be applied at the beginning of a program to prioritize radiation testing and analysis efforts in a way that most effectively minimizes risk and to provide rigorous cost justifications of these efforts based on mutually agreeable criteria.

7.0 References

1. K. A. LaBel, *et al.*, "Anatomy of an In-Flight Anomaly: Investigation of Proton-Induced SEE Test Results for Stacked IBM DRAMs", IEEE Trans. Nucl. Sci. **45**, No. 6, December 1998.
2. R. L. Pease, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 2004.
3. MIL-PRF-38535E, "General Specification for Integrated Circuits Manufacturing."
4. A Namenson and I. Arimura, "Hardness Assurance Methodology for Semiconductor Devices", IEEE Trans. Nucl. Sci. **NS-30**, No. 6, December 1983.
5..A. Namenson, "Lot Uniformity and Small Sample Sizes in Hardness Assurance", IEEE Trans. Nucl. Sci. **35**, No. 6, December 1988.
6 MIL-HDBK 814, "Ionizing Dose and Neutron Hardness Assurance Guidelines for Microcircuits and Semiconductor Devices".
7. P. McNulty, "Predicting Single-Event Phenomena in Space," Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1990.
8. A. Akkerman and J. Barak, "Ion-Track Structure and its Effects in Small Size Volumes of Silicon", IEEE Trans. Nucl. Sci. **49**, No. 6, December 2002, p. 3022.
9. R. Pease et al., "Comparison of SETs in bipolar linear circuits generated with an ion microbeam, laser light, and circuit simulation," IEEE Trans. Nucl. Sci. **49**, No. 6, December 2002, p. 3002.
10.A. Kobayashi, A. Sternberg, L. Massengill, R. Schrimpf, R. Weller, "Spatial and Temporal Characteristics of Energy Deposition by Protons and Alpha Particles in Silicon," to be published in IEEE Trans. Nucl. Sci., December 2004.
11. E. Peterson,
12. J. Barak, R. A. Reed, K. LaBel, "On the Figure of Merit Model for SEU rate calculations," IEEE Trans. Nucl. Sci. **46**, No. 6, December 1999.
13. R. Reed et al.,"Evidence for Angular Effects in Proton-Induced Single-Event Upsets," IEEE Trans. Nucl. Sci. **49**, No. 6, December 2002, p.3038.
14. S. Buchner, A. Campbell, R. Reed, P. Marshall, "Angular Dependence of Multiple-Bit Upsets Induced by Protons in a 16 Mbit DRAM," to be published in IEEE Trans. Nucl. Sci., December 2004.
15. Information on statistical distributions and inference is available at www.resacorp.com/index.htm and at www.6sigma.us/handbook/index1.html .
16. Statistical Inference.
17. Bachman, The Materials Science of Electronics.
18. D. Emily, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1996.
19. P. Dressendorfer, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1998.
20. P. Dodd, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1999.
21. J. Schwank, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 2002.
22. J. Srour, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1988.
23. L. Massengill, "SEU Modeling and Prediction Techniques," Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1993.
24. K. Galloway, G. Johnson, "Catastrophic Single-Event Effects in the Natural Space Environment," ," Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1996.

25.L.Mendoza, “Qualification from the Buyer’s Perspective,” , Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 1995.
26. .Wafer Reference.
27.T. Oldham, Nuclear And Space Radiation Effects Conference (NSREC) Short Course, New York: IEEE, 2003.
28. D. Krieg, et al., “Hardness Assurance Implications of Bimodal Total Dose Response in a Bipolar Linear Voltage Comparator, IEEE Trans. Nucl. Sci. **46**, No. 6, December 1999, p. 1627.
29. M. R. Shaneyfelt *et al*., “Impact of passivation layers on enhanced low-dose-rate sensitivity and pre-irradiation elevated-thermal stress effects in bipolar linear ICs”, IEEE Trans. Nucl. Sci. **49**, No. 6, December 2002, p. 3171.
30. R. Rivas, A. Johnston, B. Rax, T. Miyahira, M. Wiedman, “Test Results for Total Ionizing Dose Conducted at Jet Propulsion Laboratory,” to be published in Nuclear and Space Radiation Effects Conference (NSREC) Data Workshop, December 2004.
31J. Gorelick, R. Ladbury, Lina Kinchawa, “,” to be published in IEEE Trans. Nucl. Sci., December 2004.
32. J. Gorelick, R. Ladbury, Lina Kinchawa, “,”,” to be published in IEEE Trans. Nucl. Sci., December 2004.

32. A. Y. Kang, P. M. Lenahan, J. F. Conley, Jr., “The Radiation Response of the High-Dielectric Constant Hafnium Oxide/Silicon System, IEEE Trans. Nucl. Sci. **49**, No. 6, December 2002.

8.0